



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/651,691	08/29/2003	Satoshi Hada	IBM.03.55.US	9214
45728	7590	08/20/2007		
SAWYER LAW GROUP LLP P.O. BOX 51418 PALO ALTO, CA 94303			EXAMINER GORTAYO, DANGELINO N	
			ART UNIT 2168	PAPER NUMBER
			MAIL DATE 08/20/2007	DELIVERY MODE PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

**BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES**

Application Number: 10/651,691
Filing Date: August 29, 2003
Appellant(s): HADA ET AL.

MAILED

AUG 20 2007

Technology Center 2100

Erin C. Ming
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 6/12/2007 appealing from the Office action
mailed 11/14/2006

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

A statement identifying the related appeals and interferences which will directly affect or be directly affected by or have a bearing on the decision in the pending appeal is contained in the brief.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

Art Unit: 2168

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

"A Fine Grained Access Control System for XML Documents", Damiani 5-2002

Published May 2002 in "ACM Transactions on Information

and System Security", Vol. 5, No. 2, Pages 169-202.

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

Claims 1-37 are rejected under 35 U.S.C. 102(b) as being anticipated by Damiani et al. ("A Fine Grained Access Control System for XML Documents", Published May 2002 in "ACM Transactions on Information and System Security", Vol. 5, No. 2, Pages 169-202).

As per claim 1, Damiani teaches “A method for controlling access to structured documents” (see Introduction, pg. 171)

“a) providing an access control policy for a structured document comprising a plurality of nodes, wherein the access control policy comprises a plurality of access control rules;” (pg. 183, section 5.1 “Basic Features of the Access Authorizations”, wherein access authorization rules determine whether a user has access to objects)

“b) generating a path for each of the plurality of nodes in the structured document;” (pg. 174, Example 2.1 and Figure 1(a), wherein the DTD of an XML document shows path information)

“and c) generating value expression for each path based on at least one of the plurality of access control rules,” (pg. 186, Section 5.2 “Access Authorization” and Figure 5, wherein access authorizations express the requirement of access for each path of the object) “wherein the value expression is an executable statement utilized during access control evaluation to determine whether a user is allowed to access a node in the structured document.” (pg. 186, Example 5.1, Figure 5, and Algorithm 6, wherein the “Sign” column indicates the authorization for objects, as indicated by a path expression, that a user holds, as indicated in the subject column. A user is given authorization after Algorithm 6 is executed, determining the view returned to a given user accessing an object)

As per claim 2, Damiani teaches “the value expression indicates who is granted or denied access to the corresponding path associated with the node.” (pg. 186, Example 5.1 and Figure 5, wherein the “Sign” column of the access authorization table

Art Unit: 2168

indicates the subjects who are granted or denied access to each path expression associated with an object)

As per claim 3, Damiani teaches “(d) storing each path and the corresponding value expression in a table.” (pg. 186, Figure 5, wherein the access authorizations are kept in a table)

As per claim 4, Damiani teaches “(e) compiling each value expression prior to storing step (d)” (pg. 186, Example 5.1, wherein each access authorization is compiled and collected prior to placement in the table)

As per claim 5, Damiani teaches “(f) receiving a query from a user, wherein the query requests access to a node in the document;” (pg. 192, Example 6.1 lines 1-4, wherein a query from a user is received) “(g) executing the query;” (pg. 192, Example 6.1 lines 6-8, wherein the query is executed) “(h) evaluating the value expression corresponding to the path associated with the requested node;” (pg. 187, section 6.1 “Document Tree Labeling” and Figure 8, wherein the requested object’s access authorization is examined and evaluated compared to the user id) “(i) displaying data associated with the requested node if the value expression grants access to the user;” (pg. 192, Example 6.1 lines 14-21 and Figure 9(a) and 9(b), wherein the data is displayed showing accessible objects) “and (j) hiding data associated with the requested node if the value expression denies access to the user.” (pg. 192, Example 6.1 lines 14-21 and Figure 9(a) and 9(b), wherein the data is displayed hiding denied objects)

As per claim 6, Damiani teaches “the evaluating step (h) is performed during a run time.” (pg. 188, section 6.1 “Document Tree Labeling”, wherein the authorizations’ behavior varies from different requesters at runtime)

As per claim 7, Damiani teaches “wherein generating step (c) further comprises: (c1) normalizing each of the access control rules into a format comprising a head, a path and a condition, wherein the condition indicates who is granted or denied access to the path and under what circumstances;” (pg. 186, Example 5.1 and Figure 5, wherein the access authorization includes a subject, a path expression and a sign that indicated the condition) “(c2) propagating each of the plurality of access control rules through each path such that access to each path is defined by at least one access control rule;” (pg. 183, section 5.1 “Basic Features of the Access Authorizations” paragraph 2, wherein the authorizations can be recursive, propagating through the paths) “and (c3) transforming each of the at least one access control rules affecting each path into a statement indicating who is granted and denied access to the path. (pg. 183, section 5.1 “Basic Features of the Access Authorizations” paragraph 3, wherein the authorizations are indicative of who is granted or denied access, including groups)

As per claim 8, Damiani teaches “(e) replacing the value expression for a path associated with a node with a reference notation if the value expression is identical to that for a path associated with the node's parent, thereby eliminating repeated value expressions in the table.” (pg. 183, section 5.1 “Basic Features of the Access Authorizations” paragraph 2 lines 9-13, wherein recursive propagation of the authorizations applies to all descendant objects until overridden by a conflicting sign)

As per claim 9, Damiani teaches “the providing step (a) comprises: (a1) writing the plurality of access control rules; and (a2) validating the plurality of access control rules such that the resulting rules are syntactically and logically valid.” (pg. 180, section 4 “Authorization Objects”, wherein the authorizations are written and validated)

As per claim 10, Damiani teaches “the structured document is written in Extensible Markup Language. (pg. 176 paragraph 2 and Figures 1-2, wherein documents are in XML format)

As per claim 11, Damiani teaches “A computer readable medium encoded with a computer program for controlling access to a structured document” (see Introduction, pg. 171). For the remaining steps of this claim applicant(s) is/are directed to the remarks and discussions made in claim 1 above.

As per claims 12-20, these claims teach the limitations covering the same grounds as rejected claims 2-10, as discussed above, and are similarly rejected.

As per claim 21, Damiani teaches “A computer system for controlling access to a structured document,” (see Introduction, pg. 171)

“a database management system implemented on the computer system, the database management system comprising” (pg. 199, section 8.3 “The Java Implementation”)

“an access control policy for a structured document, wherein the structured document comprises a plurality of nodes and the access control policy comprises a

Art Unit: 2168

plurality of access control rules,” (pg. 183, section 5.1 “Basic Features of the Access Authorizations”, wherein access authorization rules determine whether a user has access to objects)

“and an access control mechanism configured to: generate a path for each of the plurality of nodes in the structured document” (pg. 174, Example 2.1 and Figure 1(a), wherein the DTD of an XML document shows path information)

“and generate a value expression for each path based on at least one of the plurality of access control rules,” (pg. 186, Section 5.2 “Access Authorization” and Figure 5, wherein access authorizations express the requirement of access for each path of the object)

“wherein the value expression is an executable statement utilized by the database management system during access control evaluation to determine whether a user is allowed to access a node in the structured document.” (pg. 186, Example 5.1, Figure 5, and Algorithm 6, wherein the “Sign” column indicates the authorization for objects, as indicated by a path expression, that a user holds, as indicated in the subject column. A user is given authorization after Algorithm 6 is executed, determining the view returned to a given user accessing an object)

As per claim 22, Damiani teaches “the value expression indicates who is granted or denied access to the corresponding path associated with the node.” (pg. 186, Example 5.1 and Figure 5, wherein the “Sign” column of the access authorization table indicates the subjects who are granted or denied access to each path expression associated with an object)

As per claim 23, Damiani teaches “ the Access Control mechanism is configured to store each path and the corresponding value expression in a table.” (pg. 186, Figure 5, wherein the access authorizations are kept in a table)

As per claim 24, Damiani teaches “a compiler configured to compile each value expression prior to storage of the value expression in the table.” (pg. 186, Example 5.1, and Algorithm 6, wherein each access authorization is compiled and collected prior to placement in the table)

As per claim 25, Damiani teaches “the database management system is configured to receive a query from a user, wherein the query requests access to a node in the document,” (pg. 192, Example 6.1 lines 1-4, wherein a query from a user is received) “to execute the query,” (pg. 192, Example 6.1 lines 6-8, wherein the query is executed) “to evaluate the value expression corresponding to the path associated with the requested node,” (pg. 187, section 6.1 “Document Tree Labeling” and Figure 8, wherein the requested object’s access authorization is examined and evaluated compared to the user id) “to display data associated with the requested node if the value expression grants access to the user,” (pg. 192, Example 6.1 lines 14-21 and Figure 9(a) and 9(b), wherein the data is displayed showing accessible objects) “and to hide data associated with the requested node if the value expression denies access to the user.” (pg. 192, Example 6.1 lines 14-21 and Figure 9(a) and 9(b), wherein the data is displayed hiding the denied objects)

As per claim 26, Damiani teaches “access control evaluation is performed during a run time.” (pg. 188, section 6.1 “Document Tree Labeling”, wherein the authorizations’ behavior varies from different requesters at runtime)

As per claim 27, Damiani teaches “a translator for normalizing each of the access control rules into a format comprising a head, a path and a condition, wherein the condition indicates who is granted or denied access to the path,” (pg. 186, Example 5.1 and Figure 5, wherein the access authorization includes a subject, a path expression and a sign that indicated the condition) “and for propagating each of the plurality of access control rules through each path such that access to each path is defined by at least one access control rule;” (pg. 183, section 5.1 “Basic Features of the Access Authorizations” paragraph 2, wherein the authorizations can be recursive, propagating through the paths) “and a value expression generator for transforming each of the at least one access control rules associated with each path into a statement indicating who is granted and denied access to the path.” (pg. 183, section 5.1 “Basic Features of the Access Authorizations” paragraph 3, wherein the authorizations are indicative of who is granted or denied access, including groups)

As per claim 28, Damiani teaches “the access control rules are syntactically and logically valid.” (pg. 180, section 4 “Authorization Objects”, wherein the authorizations use a standard language, XPath, for validation)

As per claim 29, Damiani teaches “the structured document is written in Extensible Markup Language.” (pg. 176 paragraph 2 and Figures 1-2, wherein documents are in XML format)

As per claim 30, Damiani teaches “A method for controlling access to structured documents” (see Introduction, pg. 171)

“a) providing an access control policy for a structured document comprising a plurality of nodes, wherein the access control policy comprises a plurality of access control rules;” (pg. 183, section 5.1 “Basic Features of the Access Authorizations”, wherein access authorization rules determine whether a user has access to objects)

“b) generating a path for each of the plurality of nodes in the structured document;” (pg. 174, Example 2.1 and Figure 1(a), wherein the DTD of an XML document shows path information)

“and c) generating value expression for each path based on at least one of the plurality of access control rules,” (pg. 186, Section 5.2 “Access Authorization” and Figure 5, wherein access authorizations express the requirement of access for each path of the object) “wherein the value expression is an executable statement utilized during access control evaluation to determine whether a user is allowed to access a node in the structured document.” (pg. 186, Example 5.1, Figure 5, and Algorithm 6, wherein the “Sign” column indicates the authorization for objects, as indicated by a path expression, that a user holds, as indicated in the subject column. A user is given authorization after Algorithm 6 is executed, determining the view returned to a given user accessing an object)

“and (d) storing each path and the corresponding value expression in a table;” (pg. 186, Figure 5, wherein the access authorizations are kept in a table) “wherein the

corresponding value expression is utilized during access control evaluation to determine whether a user is allowed to access a node in the structured document.” (pg. 186, Example 5.1, Figure 5, Algorithm 6, wherein the “Sign” column indicates the subjects who are granted access to each path expression associated with an object, used in the evaluation of views directed at a user)

As per claim 31, Damiani teaches “(e) receiving a query from a user, wherein the query requests access to a node in the document;” (pg. 192, Example 6.1 lines 1-4, wherein a query from a user is received)

“(f) executing the query;” (pg. 192, Example 6.1 lines 6-8, wherein the query is executed)

“(g) evaluating the value expression corresponding to the path associated with the requested node during a run time;” (pg. 187, section 6.1 “Document Tree Labeling” and Figure 8, wherein the requested object’s access authorization is examined and evaluated compared to the user id)

“(h) displaying data associated with the requested node if the value expression grants access to the user;” (pg. 192, Example 6.1 lines 14-21 and Figure 9(a) and 9(b), wherein the data is displayed showing accessible objects)

“and (i) hiding data associated with the requested node if the value expression denies access to the user.” (pg. 192, Example 6.1 lines 14-21 and Figure 9(a) and 9(b), wherein the data is displayed hiding denied objects)

As per claim 32, Damiani teaches “generating step (c) further comprises: (c1) normalizing each of the access control rules into a format comprising a head, a path

and a condition, wherein the condition indicates who is granted or denied access to the path and under what circumstances;" (pg. 186, Example 5.1 and Figure 5, wherein the access authorization includes a subject, a path expression and a sign that indicated the condition)

"(c2) propagating each of the plurality of access control rules through each path such that access to each path is defined by at least one access control rule;" (pg. 183, section 5.1 "Basic Features of the Access Authorizations" paragraph 2, wherein the authorizations can be recursive, propagating through the paths)

"and (c3) transforming each of the at least one access control rules affecting each path into a statement indicating who is granted and denied access to the path." (pg. 183, section 5.1 "Basic Features of the Access Authorizations" paragraph 3, wherein the authorizations are indicative of who is granted or denied access, including groups)

As per claim 33, Damiani teaches "A computer readable medium containing programming instructions for providing path-level access control to a structured document in a collection stored in a database, wherein the structured document comprises a plurality of nodes," (see Introduction, pg. 171). For the remaining steps of this claim applicant(s) is/are directed to the remarks and discussions made in claim 30 above.

As per claims 34-35, these claims teach the limitations covering the same grounds as rejected claims 31-32, as discussed above, and are similarly rejected.

As per claim 36, Damiani teaches “A method for controlling access to structured documents” (see Introduction, pg. 171)

“a) providing an access control policy for a structured document comprising a plurality of nodes, wherein the access control policy comprises a plurality of access control rules;” (pg. 183, section 5.1 “Basic Features of the Access Authorizations”, wherein access authorization rules determine whether a user has access to objects)

“b) generating a path for each of the plurality of nodes in the structured document;” (pg. 174, Example 2.1 and Figure 1(a), wherein the DTD of an XML document shows path information)

“and c) generating value expression for each path based on at least one of the plurality of access control rules,” (pg. 186, Section 5.2 “Access Authorization” and Figure 5, wherein access authorizations express the requirement of access for each path of the object)

“wherein the generating step comprising: (c1) normalizing each of the access control rules into a format comprising a head, a path and a condition, wherein the condition indicates who is granted or denied access to the path and under what circumstances;” (pg. 186, Example 5.1 and Figure 5, wherein the access authorization includes a subject, a path expression and a sign that indicated the condition) “(c2) propagating each of the plurality of access control rules through each path such that access to each path is defined by at least one access control rule;” (pg. 183, section 5.1 “Basic Features of the Access Authorizations” paragraph 2, wherein the authorizations

can be recursive, propagating through the paths) “and (c3) transforming each of the at least one access control rules affecting each path into a statement indicating who is granted and denied access to the path,” (pg. 183, section 5.1 “Basic Features of the Access Authorizations” paragraph 3, wherein the authorizations are indicative of who is granted or denied access, including groups)

“and (d) storing each path and the corresponding value expression in a table;” (pg. 186, Figure 5, wherein the access authorizations are kept in a table) “wherein the corresponding value expression is utilized during access control evaluation to determine whether a user is allowed to access a node in the structured document.” (pg. 186, Example 5.1, Figure 5, Algorithm 6, wherein the “Sign” column indicates the subjects who are granted access to each path expression associated with an object, used in the evaluation of views directed at a user)

“wherein the value expression is an executable statement utilized during access control evaluation to determine whether a user is allowed to access a node in the structured document.” (pg. 186, Example 5.1, Figure 5, and Algorithm 6, wherein the “Sign” column indicates the authorization for objects, as indicated by a path expression, that a user holds, as indicated in the subject column. A user is given authorization after Algorithm 6 is executed, determining the view returned to a given user accessing an object)

As per claim 37, Damiani teaches “A computer readable medium containing programming instructions for providing path-level access control to a structured

document in a collection stored in a database, wherein the structured document comprises a plurality of nodes" (see Introduction, pg. 171). For the remaining steps of this claim applicant(s) is/are directed to the remarks and discussions made in claim 36 above.

(10) Response to Argument

With respect to the outstanding 35 U.S.C. 102(b) rejections relating to all the independent claims, and the remaining claims which depend therefrom, Applicants argue that Damiani ("A Fine Grained Access Control System for XML Documents") does not teach "generating value expression for each path based on at least one of the plurality of access control rules, wherein the value expression is an executable statement utilized during access control evaluation to determine whether a user is allowed to access a node in the structured document" because "access authorization" as taught by Damiani in pg. 186, Section 5.2 (pg. 185) and Figure 5 is not the same thing as a value expression for each path and that the "compute-view algorithm" as shown in Figure 6 of Damiani cannot be construed as disclosing value expression is an executable statement utilized during access control evaluation.

The examiner respectfully disagrees with appellant's arguments. The examiner respectfully submits that Damiani does in fact teach generating a value expression for

Art Unit: 2168

each path based on at least one of the plurality of access control rules in Section 5.2

"Access Authorizations". The first paragraph of section 5.2 specifically states that:

"At each server, a set Auth of access authorizations specifies the actions that subjects are allowed (or forbidden) to exercise on the objects stored at the server... An access authorization $a \in \text{Auth}$ is a five-tuple of the form: {subject, object, action, sign, type}"

This is interpreted as the access authorization of Damiani is a statement defining the access rules to an object by requestors. Damiani defines an object as either a Uniform Resource Identifier (URI) of an object that defines the identity or name of an object, or of a combination of Uniform Resource Identifier and the path expression of an object, as stated in lines 6-7 of Section 5.2 (pg. 185). Figure 5 shows an example of access authorizations, where one can see that path expressions of an object is matched with the sign to determine access control.

The argument put forth by the applicant in regards to the argument that the "access authorization" of Damiani teaches "value expression" is:

"construing the "+" or "-" signs under the Sign column in the "access authorizations" table shown Figure 5 of Damiani as disclosing the "value expression" recited in claim 1. The Examiner then appears to be construing the "Compute-view algorithm" shown in Figure 6 of Damiani as disclosing the "value expression" recited in claim 1."

However, the Examiner is interpreting the "access authorization" of Damiani as teaching the "value expression" recited in claim 1. The "+" or "-" signs under the "Sign" column is read, as per Figure 5 as a part of an access authorization statement. Page 181 lines 22-42 of Damiani defines the path expression used with a Uniform Resource Identifier within the subject tuple of an access authorization statement as defining object names and gives the ability to specify individual nodes of a document.

Path expressions may also include functions. These functions serve various needs, such as the extraction of the attributes of an element and the navigation in the document structure. Table I illustrates the XPath predefined functions, their arguments type, and a brief description of the function. The name of a function and its arguments are separated by a double colon "::". For instance, expression `research/ancestor::department` returns the department nodes that appear as ancestors of the research node; expression `ps/attribute::xlink:href` returns attribute `xlink:href` of ps elements; expression `/department/child::medical staff//city` returns all the city nodes descendants of the medical staff node child of the department node (with reference to the document in Figure 2, the expression identifies the cities "Emeryville" and "Oakland"). Note that the operators *dot*, *double dot*, and *double slash* previously listed can be used as abbreviations for functions *self*, *parent*, and *descendant*, respectively. Analogously, character `@` is used as an abbreviation for function *attribute*. For instance, expression `ps/@xlink:href` is short for path expression `ps/attribute::xlink:href`. The syntax for XPath expressions also permits us to associate conditions with the nodes of a path; in this case the path expression identifies the set of nodes that satisfy all the conditions. Conditions greatly enrich the power of the language, and are a fundamental component in the construction of a sophisticated authorization mechanism.

As stated in the cited text of Damiani, the path expression can include functions that can be used for various purposes. The path expression, as coupled with the sign functions and the other tuples of an access authorization statement, are used to determine the access rights of a subject for various objects through the compute-view algorithm shown in Figure 6.

As opposed to the argument set forth by the Applicant, the "compute-view algorithm" of Damiani is not being interpreted as the "value expression" but rather to the "access control evaluation" process disclosed in claim 1 of the instant application. The "compute-view algorithm" is a process to evaluate whether a user has access permission to view a document or parts of a document.

The view of a subject on each document depends on the access permissions and denials specified by the authorizations and their priorities. Such a view can be computed through a *tree labeling process*, described next. In the following, we

use the term node (of a document tree) to refer to either an element or an attribute in the document indiscriminately.

As stated in the above cited section, the view is accomplished through a labeling process, wherein the access authorizations contained access authorization table, established earlier as analogous to "value expressions", are used in the "compute-view algorithm" to evaluate whether a user has permission to access and view an element, or node, of a document. As outlined in Section 6.1 "Document Tree Labeling" of Damiani, the authorization is read from the sign section of access authorization, and determines whether a subject, also in the access authorization, can access the object, which is also in the access authorization outlined previously and contains path expressions.

Particular attention is directed towards the InitialLabel function of the compute-view algorithm found in Figure 6 and described in page 188 lines 14-28, wherein the InitialLabel uses the access authorization, composed of a five-tuple expression containing subject, object, action, sign, and type, to determine the nodes identified in the object portion of access authorization. Once nodes are identified, the SetLabel function of the compute-view algorithm associates the objects, in this case the users, with either Allowed or Denied, permission to access objects, in this case nodes of a document, in a list (page 190 lines 5-24).

Once the algorithm outlined in Figure 6 is completed, subjects, the users of the system, are provided with permission rules for nodes in a system. A subject, whether individually or in groups, are either allowed or denied access and views to objects within a system as specified in the "sign" section of access authorizations.

The authorization model supports authorizations at all levels of granularity, including individual documents and elements within them. The object granularity for which authorizations can be specified spans from the DTD (meaning the set of its instances) to single elements/attributes within individual documents, where elements and attributes can be referenced by means of path expressions as illustrated in Section 4. Authorizations can be either positive (permissions) or negative (denials). (page 183 lines 7-13)

As summarized above, Damiani discloses value expression for each path based on at least one of the plurality of access control rules, wherein the value expression is an executable statement utilized during access control evaluation to determine whether a user is allowed to access a node in the structured document. Damiani provides a method to control access to a plurality of nodes for either individuals and groups of users. The nodes can represent either elements within documents, or attributes within documents. The access authorizations of Damiani, which includes subjects, objects, and a sign value, is utilized within a compute-view algorithm to determine user permissions to access the nodes of a document.

Conclusion:

It is respectfully submitted that the reference cited discloses the claimed data pipeline architecture that contains plural modules in series with each other to sequentially process data. In light of the forgoing arguments, the examiner respectfully requests the honorable board of Appeals and Interferences to sustain the rejection.

Art Unit: 2168.

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

Respectfully submitted,

/Dangelino Gortayo/


Dangelino Gortayo, Assistant Examiner, AU 2168

8/9/2007

Conferees:


Tim Vo, SPE AU 2168

Eddie Le, TQAS/Appeals Specialist, TC 2100



TIM VO

SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100



EDDIE C. LEE
SUPERVISORY PATENT EXAMINER

A Fine-Grained Access Control System for XML Documents

ERNESTO DAMIANI

Università di Milano

SABRINA DE CAPITANI DI VIMERCATI

Università di Brescia

STEFANO PARABOSCHI

Politecnico di Milano

and

PIERANGELA SAMARATI

Università di Milano

Web-based applications greatly increase information availability and ease of access, which is optimal for public information. The distribution and sharing of information via the Web that must be accessed in a selective way, such as electronic commerce transactions, require the definition and enforcement of security controls, ensuring that information will be accessible only to authorized entities. Different approaches have been proposed that address the problem of protecting information in a Web system. However, these approaches typically operate at the file-system level, independently of the data that have to be protected from unauthorized accesses. Part of this problem is due to the limitations of HTML, historically used to design Web documents. The extensible markup language (XML), a markup language promoted by the World Wide Web Consortium (W3C), is de facto the standard language for the exchange of information on the Internet and represents an important opportunity to provide fine-grained access control. We present an access control model to protect information distributed on the Web that, by exploiting XML's own capabilities, allows the definition and enforcement of access restrictions directly on the structure and content of the documents. We present a language for the specification of access restrictions, which uses standard notations and concepts, together with a description of a system architecture for access control enforcement based on existing technology. The result is a flexible and powerful security system offering a simple integration with current solutions.

Categories and Subject Descriptors: H.2.0 [Database Management]: General—*security, integrity, and protection*; H.2.7 [Database Management]: Database Administration—*security, integrity,*

The work reported in this article was partially supported by the Italian MURST DATA-X project and by the European Community within the Fifth (EC) Framework Programme under contract IST-1999-11791—FASTER project.

Authors' addresses: E. Damiani and P. Samarati, Dipartimento di Tecnologie dell'Informazione, Università di Milano, Via Bramante, 65, 26013 Crema (CR), Italy; email: {damiani,samarati}@dti.unimi.it; S. De Capitani di Vimercati, Dipartimento di Elettronica per l'Automazione, Università di Brescia, Via Branze 38, 25123 Brescia, Italy; email: decapita@ing.unibs.it; S. Paraboschi, Dipartimento di Elettronica e Informazione, Politecnico di Milano, Piazza L. da Vinci 32, 20133 Milano, Italy; email: parabosc@elet.polimi.it.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2002 ACM 1094-9224/02/0500-0169 \$5.00

and protection; K.6.5 [Management of Computing and Information Systems]: Security and Protection

General Terms: Security

Additional Key Words and Phrases: Access control, authorizations specification and enforcement, World Wide Web, XML documents

1. INTRODUCTION

An ever-increasing amount of information is being made available in unstructured and semistructured form via Web sites both on corporate Intranets and on the global Internet. Web-based information interchange is particularly important in electronic commerce (EC) applications, where basic transactions such as vendor registration, bidding submissions, requests for quotes, and contracts are increasingly realized by exchanging the appropriate digital documents. The huge success of the Web as a platform for EC and information dissemination has brought an increasing awareness of the fact that document exchange on the Internet should meet precise security requirements such as fine-grained authenticity, secrecy, nonrepudiation, and access control, involving data units at the level of granularity stipulated by the communicating parties. However, fully meeting these requirements through HTML-based information processing turns out to be rather awkward, due to HTML's inherent limitations. HTML provides no clean separation between the structure and the layout of a document and some of its content is only used to specify the document layout. Moreover, site designers often prepare HTML pages according to the needs of a particular browser. Therefore, HTML markup has generally little to do with data semantics. In the last few years, this situation has improved dramatically, following the standardization effort by the World Wide Web Consortium (W3C) that gave birth to the extensible markup language (XML) [Bray et al. 2000]. XML is a markup metalanguage providing semantics-aware markup without losing the formatting and rendering capabilities of HTML. XML's tags' capability of self-description is shifting the focus of Web communication from conventional hypertext to data interchange. Although HTML was defined using only a small and basic part of SGML (standard generalized markup language: ISO 8879), XML is a sophisticated subset of SGML, designed to describe data using arbitrary tags. As its name implies, extensibility is a key feature of XML; users and applications are free to declare and use their own tags and attributes. Therefore, XML ensures that both the logical structure and content of semantically rich information is retained. XML focuses on the description of information structure and content as opposed to its presentation. Presentation issues are addressed by a separate language, XSL (XML style language) [Adler et al. 2001], which is also a W3C standard for expressing how XML-based data should be rendered. In addition to XML and XSL, XLink (XML linking language), which is a specification language to define anchors and links within XML documents, is in the process of standardization [DeRose et al. 2001]. Due to its advantages, XML is now widely accepted in the Web community, and available applications exploiting this standard include OFX (open financial exchange) [CheckFree Corp 2001]

to describe financial transactions, CDF (channel definition format) [Ellerman 1997] for push technologies, and OSD (open software distribution) [van Hoff et al. 1997] for software distribution on the Net. This wealth of applications suggests that XML has great potential as an exchange format for semistructured data. The application to XML data of the latest advancement of public key cryptography has remedied most of the security problems in communication; commercial products are becoming available (such as AlphaWorks' [2001] XML Security Suite) providing fine-grained security features such as digital signatures and elementwise encryption to transactions involving XML data as a way to meet authenticity, secrecy, and nonrepudiation requirements in XML-based transactions.

The objective of our work is to complete this picture, exploiting XML's own capabilities to define and implement an authorization model for regulating access to XML documents. The rationale for our approach is defining an XML markup for a set of security elements describing the protection requirements of XML documents. Our security markup can be used to provide both instance and schema-level authorizations at the granularity of XML elements. Taken together with a user's identification and associated group memberships, as well as with the support for both permissions and denials of access, our security markup allows us to easily express different protection requirements with the support of exceptions. The enforcement of the requirements stated by the authorizations produces a view of the documents for each requester; the view includes only the information that the requester is entitled to see. A main feature of our model is that it smoothly accommodates the needs of both organizationwide policy managers and single document authors, automatically taking both into account to define who can exercise which access privileges on a particular XML document. Our notion of subject comprises identity and location; identity can include information about group or organization membership. The granularity of objects may be as fine as single elements or even attributes within XML documents. Our model includes data-dependent conditions and is open and extendable so that other enforcement conditions, such as temporal ones, could be easily added. We also present an algorithm that ensures fast online computation of such a view on XML documents requested via an HTTP connection. The proposed approach, although powerful enough to define sophisticated access to XML data, makes the design of a server-side security processor for XML documents rather straightforward. We also describe the major aspects of our Java-based implementation of the system.

1.1 Related Work

Although several projects for supporting authorization-based access control on the Web have recently been carried out, authorizations and access control mechanisms available today are at a preliminary stage. For instance, the Apache server (www.apache.org) allows the specification of access control lists via a configuration file (`access.conf`) containing the list of users, hosts (IP addresses), or host/user pairs, which must be allowed/forbidden connection to the server. Users are identified by user- and group-names and passwords, to be specified via

UNIX-style password files. By specifying a different configuration file for each directory, it is possible to define authorizations on a directory basis; files belonging to the same directory are subject to the same authorizations. Although Apache 1.2 and later also allow the protection of individual files, it is not possible to specify authorizations on portions of files. This limitation may force protection requirements to affect data organization at the file-system level. For instance, a file containing data with different protection requirements will have to be split into two different files. The proposal in Samarati et al. [1996] overcomes this limitation by allowing the specification of authorizations on portions of an HTML document. However, again, no semantic context similar to that provided by XML can be supported and the model remains limited. Some approaches, such as the EIT SHTTP scheme [Rescorla and Schiffman 1999], explicitly represent authorizations within the documents by using security-related HTML tagging. Every document may have associated security (meta)tags describing the authorizations on the document. This seems to be the right direction towards the construction of a more powerful access control mechanism, however, due to HTML fundamental limitations these proposals cannot take the information structure and semantics into full consideration.

The development of XML represents an important opportunity to solve this problem. Proposals are under development by industry and academia, and commercial products are becoming available that provide security features around XML. However, some of these approaches focus on lower-level features, such as encryption and digital signatures [AlphaWorks 2001; Eastlake et al. 2001], on query response authentication [Devanbu et al. 2001], or on privacy restrictions on the dissemination of information collected by the server [Reagle and Cranor 1999].

Work closest to ours is represented by proposals related to the specification and enforcement of security restrictions on XML documents or using XML. We first proposed the notion of a fine-grained access control model for XML documents in Damiani et al. [2000a,b], where we introduced the use of an authorization sheet associated with each XML document/DTD expressing the authorizations on the document. The approach exploiting XML's own capability as each authorization sheet is itself an XML document. In this article we extend these proposals by enriching the authorization types supported by the model, providing a complete description of the specification and enforcement mechanism, and reporting on its implementation. Among comparable proposals, Bertino et al. [2000, 2001] and Gabillon and Bruno [2001] subsequently proposed an access control environment for XML documents and some techniques to deal with authorization priorities and conflict resolution issues. A distinct, although related, line of research has been pursued by Kudoh et al. [2000], who proposed a fine-grained authorization specification language where authorizations are always associated with single elements in a document. Other related work concerns exploiting XML as a security specification language, which includes the current OASIS effort to define a standard XACML (extensible access control markup language) (<http://www.oasis.org>), and the XrML proposal [ContentGuard 2001], an extensible rights markup language (XrML) for describing usage restrictions on digital resources.

At the same time, the security community is proceeding towards the development of sophisticated access control models and mechanisms able to support different security requirements and multiple policies [Jajodia et al. 2001; Woo and Lam 1993]. These proposals have not been conceived for semistructured data with their flexible and volatile organization. They are often based on the use of logic languages, which are not immediately suited to the Internet context, where simplicity and easy integration with existing technology must be ensured. With the Web, the advantages brought by the use of a sophisticated language in terms of expressive power do not seem to be justified with respect to the added complexity. Our approach expresses security requirements in syntax, rather than in logic, leading to a simpler and more efficient evaluation engine. This characteristic ensures that our proposal can be smoothly integrated in an environment for XML information processing.

The use of authorization priorities with propagation and overriding, which is an important aspect of our proposal, may recall approaches made in the context of object-oriented databases, such as those of Fernandez et al. [1994], Jonscher et al. [1994], and Rabitti et al. [1991]. However, the XML data model is not object-oriented [Bray et al. 2000] and the hierarchies it considers represent part-of relationships and textual containment, which require specific techniques different from those applicable to ISA hierarchies in the object-oriented context.

1.2 Outline of the Article

The article is organized as follows. Section 2 illustrates the basic characteristics of the XML proposal. Sections 3 and 4 discuss the subjects and the objects, respectively, of our authorization model. Section 5 presents the authorizations supported by the access control model for expressing security requirements on XML documents. Section 6 introduces the process for computing, at access control time, the document view to be returned to the requester and presents an algorithm for efficiently computing such a view; it also discusses data modeling issues taking into account the schema changes due to the partial view. Section 7 discusses the support of write actions. Section 8 addresses design and implementation issues and illustrates the architecture of the access control processor. Section 9 presents our concluding remarks.

2. PRELIMINARY CONCEPTS

XML [Bray et al. 2000] is a markup language for describing semistructured information. An XML document is composed of a sequence of nested elements, each delimited either by a pair of start and end tags (e.g., <nurse> and </nurse>) or by an empty tag (e.g., <organization/>). XML documents can be classified into two categories: *well-formed* and *valid*. An XML document is well-formed if it obeys the syntax of XML (e.g., nonempty tags must be properly nested; each nonempty start tag must correspond to an end tag). A well-formed document is valid if it conforms to a proper document type definition (DTD). A DTD is a file (external, included directly in the XML document, or both) that contains a formal definition of a particular type of XML documents.

A DTD may include declarations for elements, attributes, entities, and notations. Elements are the most important components of an XML document. Element declarations in the DTD specify the names of elements and their content. The content specification may coincide with Empty, Any, or with a group of one or more subelements/groups. Empty means that the element has no content, whereas Any means that the element may have any content. Groups can be sequences or a choice of subelements and/or subgroups. Sequences of elements are separated by a comma “,” and choices are separated by the vertical bar “|”. Declarations of sequence and choices of subelements also describe the subelements’ cardinality; with a notation inspired by extended BNF grammars, “*” indicates zero or more occurrences, “+” indicates one or more occurrences, “?” indicates zero or one occurrence, and no label indicates exactly one occurrence. XML also allows us to declare elements with a *mixed* content, that is, elements containing parsable character data (#PCDATA), optionally interleaved with subelements. Attributes represent properties of elements. Attribute declarations specify the attributes of each element, indicating their name, type, and, possibly, default value. Attributes can be marked as required, implied, or fixed. Attributes marked as required must have an explicit value for each occurrence of the elements with which they are associated. Attributes marked as implied are optional. Attributes marked as fixed have a fixed value indicated at the time of their definition. Entities are used to include text and/or binary data in a document and can be *internal* or *external*. Internal entities are used to introduce special characters in the document or as an alias for some frequently used text. External entities are external files containing either text or binary data. Notation declarations specify how to manage the binary entities. Entities and notations are important in the description of the physical structure of an XML document, but are not considered in this article, where we concentrate the analysis on the XML logical description. Our authorization model can be easily extended to cover these components.

XML documents may include *links* that express relationships between resources. A link is defined by an XLink (XML linking language) *linking element* [DeRose et al. 2001]. XLink allows two types of links: *simple links*, similar to the HTML links, and *extended links*, that express relationships among more than two resources. We refer the reader to the W3C proposal [DeRose et al. 2001] for a complete description of XLink elements and attributes.

Example 2.1. Figure 1(a) illustrates an example of DTD for XML documents describing a department of a hospital. Each department is composed of zero, one, or more divisions and is responsible for creating an XML document for each division (if any) conforming to the considered DTD. According to this DTD, a department includes elements *division* (optional), *medical_staff*, *research activity*, and *patients*. The medical staff is composed of physicians and nurses (physician and nurse elements). Each physician is characterized by name, specialty, office, phone, home address, and salary. Each nurse is characterized by name, home address, and salary. The research activity of a department (or division) is organized into projects, each with a designated leader and consisting of an objective, laboratory, and a set, possibly empty,



(b)

of related publications. Each publication has a title, one or more author elements, and the corresponding postscript file (linking element `ps`). Information about the patients includes name, address, room, illness, and therapy. A therapy is characterized by its startdate and enddate, a type, and a list of drugs. Each drug prescribed for a patient has a name, a daily administration quantity (daily_admin element), and a cost. Properties of each element are defined in the attribute definition portion of the document. Elements department and project have a required attribute name, which is a string (character data). In addition, element project has an attribute type representing the project type (public vs. private). Attribute `xmlns:xlink` of element `ps` is used to define an XLink namespace,¹ `xlink:type` denotes the type of the link, and `xlink:href` is the locator attribute that defines where the resource is located.

XML documents valid according to a DTD obey the structure defined by the DTD. Figure 2 illustrates an example of an XML document valid with respect to the DTD in Figure 1. Intuitively, each DTD is a schema, and XML documents valid according to that DTD are instances of that schema. Note that since elements and attributes defined in a DTD may appear in an XML document zero (optional elements), one, or multiple times, according to their cardinality constraints, the structure specified by the DTD is not rigid; two distinct documents of the same schema may differ in the number and structure of elements.

DTDs and XML documents can be modeled graphically as follows. A DTD is represented as a labeled tree containing a node for each element, attribute, and value associated with fixed attributes in the DTD. There is an arc between an element and an element/attribute belonging to it, labeled with the cardinality of the relationship, and between a fixed attribute and each of its value(s). Figure 1(b) illustrates the tree for the DTD in Figure 1(a). Elements are represented as ovals and attributes as rectangles. Arcs labeled `or` and with multiple branching are used to represent a choice in an element declaration. For instance, choice `(city|county)?` in the address declaration is represented by an arc, labeled with “?” and “or”, from address node to both city and county nodes. An arc with multiple branching is also used to represent a sequence with a cardinality constraint associated with the whole sequence. For instance, sequence `(author+,title,ps?)*` in the publications declaration is represented with an arc labeled “*” starting from node publications and ending at nodes author, title, and ps. To preserve the order between elements in an element declaration, for any two elements e_i and e_j , if e_j follows e_i in the element declaration, node e_j appears below node e_i in the tree. Each XML document is described by a tree with a node for each element, attribute, and value in the document, and with an arc between each element and each of its subelements/attributes/values and between each attribute and each of its value(s). Each arc in the DTD tree may correspond to zero, one, or multiple arcs in the XML document, depending

¹A namespace is a way of distinguishing element types and attribute names to allow correct processing by a software module [Bray et al. 1999]. For instance, an XML document may include one or more occurrences of the element `title` to represent either a paper's title or an identifying appellation such as Mr. or Professor. The different semantics is captured by associating them with different namespaces.

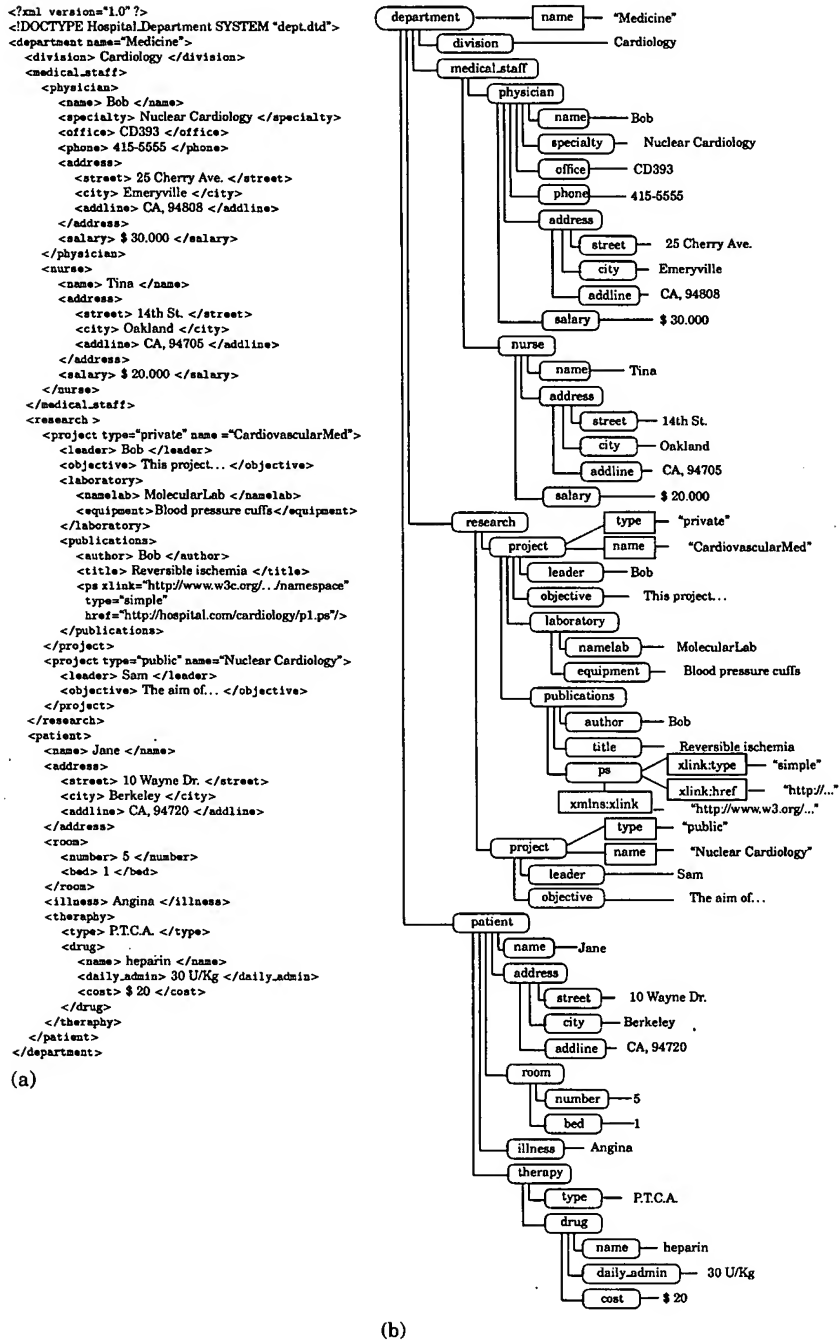


Fig. 2. (a) An example of valid XML document conforming to the DTD in Figure 1 and (b) the corresponding graphical representation.

on the cardinality of the corresponding containment relationship. Note that arcs in XML documents are not labeled. Figure 2(b) illustrates the tree representation of the XML document in Figure 2(a). In the remainder of this article we use the term “tree” to indiscriminately refer to the graphical representation of either a DTD or an XML document. Also, we use the term “object” to refer to XML documents and DTDs indiscriminately. We explicitly distinguish between them when necessary.

Each object can have associated *metadata*. These are data representing information on the object, such as creator, creation date, expiration date, as well as any other properties that may have been defined on it (e.g., public document, internal document, and so on). Metadata can be conveniently managed through a description in RDF (resource description framework) [Brickley and Guha 2000], an application of XML resulting from a collaborative design effort among several W3C members. The RDF description associated with an object can be seen as a set of pairs of the form $\langle \textit{property}, \textit{value} \rangle$, where *property* is the name of a metaproperty and *value* is its value with respect to the object. The value can be atomic (e.g., string or number), another resource, a collection of atomic values, or a metadata instance. Metaproperties can be nested and their structure represented in a graphical framework similar to that of XML and DTD documents.

3. AUTHORIZATION SUBJECTS

The development of an access control system requires the definition of the subjects and objects against which authorizations must be specified and access control must be enforced. In this section we present the subjects; in Section 4 we describe the objects.

Usually, subjects can be referred to on the basis of their identities or on the location from which requests originate. Locations can be expressed with reference to either their numeric IP address (e.g., 150.100.30.8) or their symbolic name (e.g., tweety.cardiology.hospital.com). Our model combines these features. Subjects requesting access are thus characterized by a triple (user-id, IP-address, sym-address), where user-id is the identity² with which the user connected to the server, and IP-address (sym-address, respectively) is the numeric (symbolic, respectively) identifier of the machine from which the user connected.

To allow the specification of authorizations applicable to sets of users and/or sets of machines, our model also supports user *groups* and *location patterns*. A group is a set of users defined at the server. Groups do not need to be disjoint and can be nested. A location pattern is an expression identifying a set of physical locations, with reference to either their symbolic or numerical identifiers. Patterns are specified by using the wild card character * instead of a specific name

²We assume user identities to be local, that is, established and authenticated by the server, because this is a solution relatively easy to implement securely. Obviously, in a context where remote identities cannot be forged and can therefore be trusted by the server (using a Certification Authority, a trusted third party, or any other secure infrastructure), remote identities could be considered as well.

or number (or sequence of them). For instance, $151.100.*.*$, or equivalently $151.100.*$, denotes all the machines belonging to network 151.100 . Similarly, $*.mil$, $*.com$, and $*.it$ denote all the machines in the military, company, and Italy domains, respectively. If multiple wild card characters appear in a pattern, their occurrence must be continuous (not interleaved by numbers or names). Also, consistent with the fact that specificity is left to right in IP addresses and right to left in symbolic names, wild card characters must always appear as rightmost elements in IP patterns and as leftmost elements in symbolic patterns. Intuitively, location patterns are to location addresses what groups are to users.

Users and groups together with their membership relationship, IP addresses with patterns, and symbolic names with their patterns, form partially ordered sets (hierarchies). To provide a uniform treatment for the different components of subjects, we first give the definition of hierarchy as follows.

Definition 3.1 Hierarchy. A hierarchy is a triple (X, Y, \leq) , where \leq is a partial order on Y , called the dominance relation, and $X \subseteq Y$ is the set of minimal elements of Y with respect to the partial order.

The model considers the following three hierarchies.

- A *user-group hierarchy* $UGH = (U, UG, \leq_{UG})$, where U is a set of user identifiers and $UG = U \cup G$, with G a set of group names in which users are organized. Given two elements $x, y \in UG$, $x \leq_{UG} y$ if and only if x is a member of y . In most practical applications the hierarchy is rooted at a group, called Public or Any, to whom everybody (directly or indirectly) belongs.
- An *IP hierarchy* $IPH = (I, IP, \leq_{IP})$, where I is a set of completely specified numerical addresses and IP is a set of IP patterns. Given two elements $x, y \in IP$, $x \leq_{IP} y$ if and only if each component of y is either the wild card character or is equal to the corresponding (positionwise from left to right) component of x .
- A *symbolic name hierarchy* $SNH = (S, SN, \leq_{SN})$, where S is a set of completely specified symbolic names and SN is a set of symbolic name patterns. Given two elements $x, y \in SN$, $x \leq_{SN} y$ if and only if each component of y is either the wild card character or is equal to the corresponding (positionwise from right to left) component of x .

A hierarchy can be pictured as a directed graph containing a node for each element of the hierarchy and with an arc from an element x to an element y , if x directly dominates y in the hierarchy. Dominance relationships holding in the hierarchy correspond to paths in the graph. User-group hierarchies are arbitrary DAGs, and IP and symbolic name hierarchies are necessarily trees. Figure 3 illustrates an example of user-group, IP, and symbolic name hierarchies. We note that the hierarchies introduced serve explanatory purposes and only the user-group hierarchy needs to be explicitly defined and stored at the server (or retrieved at access control time).

Instead of specifying authorizations with respect to only one of either the user/group identifier or location identifier, and having the problem of how




Fig. 3. (a) An example of user-group, (b) IP, and (c) symbolic name hierarchies.

different authorizations should be combined at access request time, we allow the specification of authorizations with reference to both user/group and location. This choice provides more expressiveness (it allows us to express the same requirements as the alternative and more) and provides a natural treatment for different authorizations applicable to the same request.³ We define the *authorization subject hierarchy* as the Cartesian product of the three hierarchies previously introduced, where a subject s_j is dominated by another subject s_i if each of s_j 's components is dominated by the corresponding component in s_i , as clarified by the following definition.

Definition 3.2 Authorization subject hierarchy. The *authorization subject hierarchy* is a hierarchy $ASH = (R, AS, \leq_{AS})$, where $R = \langle U \times I \times S \rangle$, $AS = \langle UG \times IP \times SN \rangle$, and $\langle ug_i, ip_i, sn_i \rangle \leq_{AS} \langle ug_j, ip_j, sn_j \rangle$, if and only if $ug_i \leq_{UG} ug_j$, $ip_i \leq_{IP} ip_j$, and $sn_i \leq_{SN} sn_j$.

Authorizations can be defined with reference to any of the elements of ASH. In particular, authorizations can be specified for users/groups regardless of the physical location (e.g., $\langle Alice, *, * \rangle$), for physical locations regardless of the user identity (e.g., $\langle Public, 150.100.30.8, * \rangle$), or for both (e.g., $\langle Sam, *, *.acme.com \rangle$). Intuitively, authorizations specified for subject $s_j \in ASH$ are applicable to all subjects s_i such that $s_i \leq_{AS} s_j$. Possible conflicts between authorizations applicable to a given requester are investigated in Section 5.

4. AUTHORIZATION OBJECTS

A set Obj of uniform resource identifiers (URI) [Berners-Lee et al. 1998] denotes the resources to be protected. For XML documents, URIs can be extended with *path expressions*, which are used to identify the elements and attributes within a document. In particular, we adopt a W3C proposal for the identification of internal components of an XML document, namely, the XPath language [World Wide Web Consortium (W3C) 2001]. There are considerable advantages deriving from the adoption of a standard language. First, the syntax and semantics of the language are known by potential users and well studied. Second, several tools are already available that can be easily reused to produce a functioning system.

³We elaborate on this in Section 5.

We keep at a simplified level the description of the language that expresses patterns in XPath. The W3C proposal [World Wide Web Consortium (W3C) 2001] contains the complete specification of the language.

Definition 4.1 Path expression. A *path expression* on a document tree is a sequence of element names or predefined functions separated by character / (slash): $l_1/l_2/\dots/l_n$. Path expressions may terminate with an attribute name as the last term of the sequence. Attribute names are syntactically distinguished by preceding them with special character @.

A path expression $l_1/l_2/\dots/l_n$ on a document tree represents all the attributes or elements named l_n that can be reached by descending the document tree along the sequence of nodes named l_1, l_2, \dots, l_{n-1} . For instance, path expression `/department/medical_staff/physician` denotes the physician elements that are children of `medical_staff` elements, that are children of the department element. Path expressions may start from the root of the document (if the path expression starts with a slash, it is called *absolute*) or from a predefined starting point in the document (if the path expression starts directly with an element name, it is called *relative*). The path expression may also contain the operators *dot*, which represents the current node; *double dot*, which represents the parent node; and *double slash*, which represents an arbitrary descending path. For instance, path expression `/department//leader` retrieves all the elements leader descendants (at any level) of the document root department.

Path expressions may also include functions. These functions serve various needs, such as the extraction of the attributes of an element and the navigation in the document structure. Table I illustrates the XPath predefined functions, their arguments type, and a brief description of the function. The name of a function and its arguments are separated by a double colon “:.”. For instance, expression `research/ancestor::department` returns the department nodes that appear as ancestors of the research node; expression `ps/attribute::xlink:href` returns attribute `xlink:href` of `ps` elements; expression `/department/child::medical_staff//city` returns all the city nodes descendants of the `medical_staff` node child of the department node (with reference to the document in Figure 2, the expression identifies the cities “Emeryville” and “Oakland”). Note that the operators *dot*, *double dot*, and *double slash* previously listed can be used as abbreviations for functions *self*, *parent*, and *descendant*, respectively. Analogously, character @ is used as an abbreviation for function *attribute*. For instance, expression `ps/@xlink:href` is short for path expression `ps/attribute::xlink:href`. The syntax for XPath expressions also permits us to associate conditions with the nodes of a path; in this case the path expression identifies the set of nodes that satisfy all the conditions. Conditions greatly enrich the power of the language, and are a fundamental component in the construction of a sophisticated authorization mechanism. The conditional expressions used to represent conditions may operate on the “text” of elements (i.e., the character data in the elements) or on names and values of attributes. Conditions are distinguished from navigation specifications by enclosing them within square brackets. Given a path expression $l_1/\dots/l_n$ on the tree of an XML document, a condition may be defined

Table I. XPath Predefined Functions

Function	Argument	Description
ancestor	<i>element-name</i>	returns all <i>element-name</i> ancestors of the context node
ancestor-or-self	<i>element-name</i>	returns all <i>element-name</i> ancestors of the context node and, if the context node is an <i>element-name</i> element, the context node as well
descendant	<i>element-name</i>	returns all <i>element-name</i> descendants of the context node
descendant-or-self	<i>element-name</i>	returns all <i>element-name</i> descendants of the context node and, if the context node is an <i>element-name</i> element, the context node as well
following	<i>element-name</i>	returns all the <i>element-name</i> nodes that are after the context node in the document order, excluding any descendants, attribute nodes, and namespace nodes
following-sibling	<i>element-name</i>	returns all the following <i>element-name</i> siblings of the context node
preceding	<i>element-name</i>	returns all the <i>element-name</i> nodes that are before the context node in the document order, excluding any ancestors, attribute nodes, and namespace nodes
preceding-sibling	<i>element-name</i>	returns all the preceding <i>element-name</i> siblings of the context node
parent	<i>element-name</i>	returns the parent of the context node, if there is one and it is an <i>element-name</i> element, and otherwise returns nothing
child	<i>element-name</i>	returns all <i>element-name</i> elements children of the context node
self	<i>element-name</i>	returns the context node, if it is an <i>element-name</i> element, and otherwise returns nothing
attribute	<i>attribute-name</i>	returns attribute <i>attribute-name</i> of the context node
namespace	<i>namespace</i>	returns the <i>namespace</i> nodes of the context node

on any label l_i , enclosing in square brackets a separate evaluation context containing a predicate that compares the result of the evaluation of the relative path expression with a constant or another expression. Conditional expressions may be combined via and and or operators to build Boolean expressions. Multiple conditional expressions appearing in the same path expression are considered to be anded (i.e., all the conditions must be satisfied). In addition, conditional expressions may include functions `last()` and `position()` that permit the extraction of the children of a node which are in given positions. Function `last()` evaluates to true on the last child of the current node. Function `position()` evaluates to true on the node in the evaluation context whose position is equal to the context position. For instance, expression `/department/research/project/publications/ps[position()=1]` returns the first `ps` child of the `publications` element (note that the conditional expression `[position()=1]` can be abbreviated as `[1]`); expression `/department[./@name = "Medicine" and ./division = "Cardiology"]/medical_staff/nurse` identifies all the nurses of the "Cardiology" division of the "Medicine" department; expression `/department/research/project[./@type="public"][1]` returns the first "public" projects child of the `research` element.

The proposed formalism also can be used to specify conditions on metadata, such as the RDF descriptions illustrated in Section 2. In this case, a new predefined function `meta` is used to access metainformation on the XML document.

For instance, expression `/[meta()]/@creation_date = "2000-01-05"`, evaluates to true for documents created on January 5, 2000.

5. ACCESS AUTHORIZATION SPECIFICATIONS

We first describe the basic features that access authorizations need to provide to regulate access to Web documents and then we give their formal definition.

5.1 Basic Features of the Access Authorizations

The authorization model supports authorizations at all levels of granularity, including individual documents and elements within them. The object granularity for which authorizations can be specified spans from the DTD (meaning the set of its instances) to single elements/attributes within individual documents, where elements and attributes can be referenced by means of path expressions as illustrated in Section 4. Authorizations can be either *positive* (permissions) or *negative* (denials). The reason for having both positive and negative authorizations is to provide a simple and effective way to specify authorizations applicable to sets of subjects/objects with support for exceptions [Jajodia et al. 2001; Lunt 1989].

Authorizations specified on an element can be defined as applicable to the element's attributes only (*local* authorizations) or, in a recursive approach, to its subelements and their attributes (*recursive* authorizations). Local authorizations on an element apply to the direct attributes of the element but not to those of its subelements. As a complement, recursive authorizations, by propagating permissions/denials from nodes to their descendants in the tree, represent an easy way to specify authorizations holding for the whole structured content of an element (on the whole document if the element is the root). To support exceptions (e.g., the whole content, but a specific element can be read), recursive propagation from a node applies until stopped by an explicit conflicting (i.e., of different sign) authorization on the descendants. Intuitively, authorizations propagate until overridden by an authorization on a more specific object [Jajodia et al. 2001].

Authorizations can be specified on single XML documents (document- or instance-level authorizations) or on DTDs (DTD- or schema-level authorizations). Authorizations specified on a DTD are applicable (i.e., are propagated) to all XML documents that are instances of the DTD. Since large enterprises are often organized into multiple domains, protection requirements may be specified both at the level of the enterprise, stating general regulations that should hold, and at the level of specific domains (part of the enterprise) where, according to a local policy, additional constraints may need to be specified or some constraints may need to be relaxed. Organizations specify authorizations with respect to DTDs; specific sites can specify authorizations with respect to individual documents (instance-level authorizations) as well as with respect to DTDs. The two types of DTD-level authorizations have complementary roles in increasing access control flexibility. Organization DTD-level authorizations stated by a central authority can be effectively used to implement corporatewide

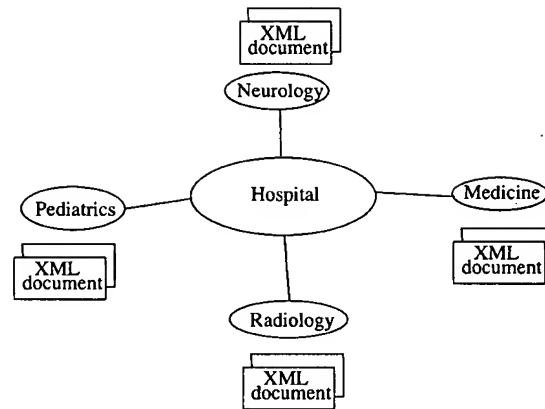


Fig. 4. An example of hospital organization.

access control policies on document classes. Site DTD-level authorizations specified by departmental authorities allow for departmentwide access control policies complementing the corporate ones. Moreover, they alleviate administration chores by allowing concise specification of sitewide authorizations. For instance, suppose that a hospital is composed of different departments, each of which is responsible for managing specific XML documents (see Figure 4). In this scenario, general protection requirements that should be satisfied by all departments of the hospital can be expressed through (organization) DTD-level authorizations stated at the hospital level. Specific protection requirements, applicable only within a single department, can be expressed by means of (site) DTD-level authorizations. Analogously, requirements applicable only to a specific document are expressed by means of instance-level authorizations associated with the document. We anticipate that, in the access control processing, organization DTD-level authorizations and site DTD-level authorizations are, with respect to each DTD, merged by performing a *flat union*. In other words, organizationwide and site-specific authorizations are treated in the same way (although, remember, organizationwide authorizations apply to all the documents in the network whereas site-specific authorizations apply only to documents stored at the site). Given this, in the following we simply refer to DTD authorizations without making any distinction of where they have been specified. The reason for merging the two sets of authorizations with a simple flat union is simplicity. We observe that, in principle, even at this level some notion of “specificity” could be applied. This reasoning could also be possibly extended by considering any number of intermediate organizational levels that could be reflected in priorities associated with the authorizations [Jonscher et al. 1994]. Authorizations at the DTD level, together with path expressions with conditions, provide an effective way for specifying authorizations on elements of different documents, possibly in a content-dependent way. Again, according to the “most specific takes precedence” principle, DTD-level authorizations being propagated to an instance are overridden by possible authorizations specified for the instance. To address situations where these precedence criteria should

Table II. Authorization Types

Level/Strength	Propagation	
	Local	Recursive
Instance	L	R
Instance (soft statement)	LS	RS
DTD	LD	RD
DTD (hard statement)	LDH	RDH

not be applied (e.g., cases where an authorization on a document should be applicable unless otherwise stated at the DTD level, or cases where an authorization on a DTD must be applied to all instances of the DTD), we allow users to specify instance-level authorizations as *soft* and DTD-level authorizations as *hard*. Nonsoft and nonhard authorizations have the behavior sketched above (i.e., nonsoft instance-level authorizations have priority over nonhard DTD-level authorizations). Soft authorizations are authorizations that apply to the document unless otherwise stated at the DTD level (intuitively, a department can state that its documents can/cannot be accessed unless the organization states otherwise). In a dual way, hard authorizations allow an organization to specify authorizations that must be enforced in all instances of a DTD, no exceptions. The combination of the options above introduces the eight authorization types summarized in Table II. Their semantics dictates a priority order among the authorization types. The priority order from the highest to the lowest is: LDH (local hard authorization), RDH (recursive hard authorization), L (local authorization), R (recursive authorization), LD (local authorization specified at the schema level), RD (recursive authorization specified at the schema level), LS (local soft authorization), and RS (recursive soft authorization). For instance, if there are a positive local hard authorization and a negative local authorization both applicable to the same object and subject, the positive local hard authorization overrides the negative one. However, it also may be the case that several authorizations, possibly of different sign, apply to a given request with reference to a given authorization type and element/attribute. As we show in the following section, conflicts between such authorizations are solved by applying a conflict resolution policy [Jajodia et al. 2001; Lunt 1989].

5.2 Access Authorizations

At each server, a set *Auth* of access authorizations specifies the actions that subjects are allowed (or forbidden) to exercise on the objects stored at the server. Access authorizations are formally defined as follows.

Definition 5.1 Access Authorization. An access authorization $a \in \text{Auth}$ is a five-tuple of the form: (subject, object, action, sign, type), where:

- *subject* $\in \text{AS}$ is the subject to whom the authorization is granted;
- *object* is either a URI in *Obj* or is of the form *URI:PE*, where *URI* $\in \text{Obj}$ and *PE* is a path expression on the tree of document *URI*;

	Subject user/group,IP,domain	Object (path expression)	Action	Sign	Type
a	Public,*,*	/department/ename	read	+	L
b	Public,*,*	/department/division	read	+	L
c	Administrative,*,*.hospital.com	/department/ename	read	+	LDH
d	Administrative,*,*.hospital.com	/department/address	read	+	RDH
e	Administrative,169.101.80.6,*	/department/medical_staff/salary	read	+	LDH
f	Administrative,169.101.80.6,*	/department/patient/cost	read	+	LDH
g	Public,*,*	/department/medical_staff/salary	read	-	LDH
h	Public,*,*	/department/patient/cost	read	-	LDH
i	Public,*,*	/department[./@name="medicine"]/medical_staff	read	+	R
j	Public,*,*	/department[./@name="medicine"]/medical_staff/address	read	-	R
k	Public,*,*	/department[./@name="medicine"]/medical_staff/salary	read	-	L
m	PhyC,*,*	/department[./@name="medicine" and ./division="cardiology"]/patient	read	+	R
n	Public,*,*	/department[./@name="medicine" and ./division="cardiology"]/patient	read	-	R
o	MedicalStaff,*,*	/department[./@name="medicine"]/research	read	+	R
p	Public,*,*	/department[./@name="medicine"]/research	read	-	R
q	PhyC,169.101.80.6,*	/department/research/project[./@type="private"]	read	+	R
r	*,*,*	/department/research/project[./@type="private"]	read	-	R
s	NurseC,*,*	/department/patient/illness	read	+	LS
t	NurseC,*,*	/department/patient/name	read	+	L
u	NurseC,*,*	/department/patient/drug	read	+	R
v	NurseC,*,*	/department/patient/room	read	+	R

Fig. 5. Example of access authorizations.

- *action* = read is the action being authorized or forbidden;⁴
- *sign* $\in \{+, -\}$ is the sign of the authorization, which can be positive (allow access) or negative (forbid access);
- *type* $\in \{LDH, RDH, L, R, LD, RD, LS, RS\}$ is the type of the authorization (Local DTD Hard, Recursive DTD Hard, Local, Recursive, Local DTD, Recursive DTD, Local Soft, and Recursive Soft, respectively).

Example 5.1. Consider the XML document in Figure 2. This document is an instance of the DTD in Figure 1 that includes information regarding the Cardiology division of the Medicine department of a given hospital. We now illustrate some examples of protection requirements that may need to be expressed and their representation as authorizations in our model. The boldface letters between square brackets at the end of each requirement identify the authorizations in Figure 5 expressing the requirement. Figure 5 lists the resulting authorizations. The horizontal line between authorizations **p** and **q** separates DTD-level authorizations (**a** through **p**) from instance-level authorizations (**q** through **v**). Note that for simplicity, in the object field we report only the path expression and omit the URI.

Hospital's policy. (Organization DTD-level authorizations applicable to all the departments of the hospital). Requirements expressed as “must” specify statements that do not allow exceptions (which translate to hard authorizations).

1. Department and division names are publicly accessible. **[a]-[b]**
2. Information about the name and home address of medical staff and of patients *must* be accessible to the members of Administrative group connected from domain *.hospital.com. **[c]-[d]**

⁴We limit our consideration to read authorizations. The support of write actions such as insert, update, and delete does not complicate the authorization model (see Section 7). However, full support for such actions in the framework of XML has yet to be defined.

3. Information about the salary of the medical staff and the cost of the therapy of all patients of the hospital *must* be accessible to the members of group Administrative connected from host 159.101.80.5. [e]-[f]
Everybody else *must* be explicitly forbidden access to this information. [g]-[h]

Medicine **department's policy.** (Site DTD-level authorizations to complement or override the organization DTD-level authorizations).

4. Information about medical staff working in the Medicine department with exception of their salary and home address, is publicly accessible. [i]-[j]-[k]
5. Information about patients hospitalized in a given division is accessible only to the physicians working in the same division. [m]-[n]
6. Information about the research activity of the Medicine's divisions is accessible only to the medical staff of the hospital. [o]-[p]

Cardiology **division's policy.** (Specified at the instance level to complement or override the hospital's and department's policy).

7. Information about "private" projects of the Cardiology division is accessible to the physicians working in the Cardiology division when connected from network 159.*. [q]
No one else can access information about "private" projects. [r]
8. Information about patients' illnesses is accessible to nurses of the Cardiology division unless otherwise stated at the DTD level. [s]
9. Information about name, drug, and room of patients hospitalized in the Cardiology division is accessible to the members of NurseC group. [t]-[u]-[v]

The following section discusses the interpretation of authorizations to produce the view of a requesting subject on the document requested.

6. REQUESTER'S VIEW ON DOCUMENTS

The view of a subject on each document depends on the access permissions and denials specified by the authorizations and their priorities. Such a view can be computed through a *tree labeling process*, described next. In the following, we use the term node (of a document tree) to refer to either an element or an attribute in the document indiscriminately.

6.1 Document Tree Labeling

Each access authorization states whether a subject can (or cannot) access an element/attribute (or set of them). The type associated with each authorization on a given object (at the instance or schema level) determines the "behavior" of the authorization with respect to the object structure, that is, whether it propagates down the tree and whether it is overridden or it overrides other authorizations. The enforcement of the authorizations on a document according to the principles discussed in Section 5.1 essentially requires

the indication of whether, for an element/attribute in a document, a positive authorization (+), a negative authorization (−), or no authorization (ϵ) applies. Since authorizations can be of different levels (instance vs. schema), strength (hard vs. soft), and propagation (local vs. recursive), we associate with each node n an array, $n.\text{veclabel}$, of eight components,⁵ one for each authorization type, which is a record including three fields: *sign*, *Allowed*, and *Denied*. Let $t \in \{\text{LDH}, \text{RDH}, \text{L}, \text{R}, \text{LD}, \text{RD}, \text{LS}, \text{RS}\}$ be an authorization type. The value of $n.\text{veclabel}[t].\text{sign}$ can be “+” for permission, “−” for denials, and “ ϵ ” for no authorization, and it indicates the sign associated with the node according to the authorizations and the conflict resolution policy. The determination of the sign is preceded by the computation of $n.\text{veclabel}[t].\text{Allowed}$ and $n.\text{veclabel}[t].\text{Denied}$, two lists storing all the subjects for which there is a positive (negative, respectively) authorization of type t that applies to n . Figure 6 illustrates an algorithm, **Compute-view**, enforcing the labeling process. Given a requester rq and an XML document URI , the algorithm first initializes variable T to the tree representing the document and r to the root of T . After initialization, the algorithm invokes procedure **InitialLabel**(T, rq). The purpose of **InitialLabel** is to associate authorizations with the corresponding elements/attributes. Since not all the authorizations defined on a document are applicable to all requesters, the set of authorizations on the document’s elements, and the authorizations’ behavior along the tree, can vary for different requesters. Thus the first step of **InitialLabel** consists of the determination of the set A of authorizations defined for the document URI at the instance and schema levels and applicable to the requester rq . For each authorization $a = \langle \text{subject}, \text{object}, \text{action}, \text{sign}, \text{type} \rangle$ in A , the method determines the set N of nodes in T that are identified by $a.\text{object}$. After that, for each node n in N , **InitialLabel** adds $a.\text{subject}$ to $n.\text{veclabel}[a.\text{type}].\text{Allowed}$ if $a.\text{sign}$ is +; it adds the subject to $n.\text{veclabel}[a.\text{type}].\text{Denied}$ if $a.\text{sign}$ is −. Since several authorizations, possibly of different sign, may exist for each authorization type t (i.e., $n.\text{veclabel}[t].\text{Allowed}$ and $n.\text{veclabel}[t].\text{Denied}$ both can be not empty), the determination of the $n.\text{veclabel}[t].\text{sign}$ value requires the application of a conflict resolution policy. Different approaches can be used to solve these conflicts [Samarati and De Capitani di Vimercati 2001]. One solution is to consider the authorization with the most specific subject (“*most specific subject takes precedence*” principle), where specificity is dictated by the partial order defined over ASH; other solutions can consider the negative authorization (“*denials take precedence*”), the positive authorization (“*permissions take precedence*”), or no authorization (“*nothing takes precedence*”). Other approaches could also be envisioned, such as, for example, considering the sign of the authorizations that are more numerous. For simplicity, in the examples and discussion in the remainder of this article we refer to a specific policy and solve conflicts with respect to the “most specific subject takes precedence” principle and, in cases where conflicts remain unsolved (the conflicting authorizations have incompatible subjects), we stay on the safe side and apply the “*denials take precedence*”

⁵We use a Java-like notation where obj. att (obj.meth , respectively) denotes the attribute (method, respectively) associated with object obj .

ALGORITHM 6.1. *Compute-view algorithm***Input:** A requester rq and an XML document URI **Output:** The document to be returned to rq

```

void main()
{
    SecureDocument T(URI) /* Constructor creates tree from the XML document URI */
    r = T.Root;
    InitialLabel(T, rq);
    r.SetLabel();
    r.GetFinallabel(empty); /* empty is a labeling vector whose components have
                             all sign fields equal to  $\epsilon$  */
    r.Prune();
}

void InitialLabel(T, rq)
{
    A = {a = (subject, object, action, sign, type) | a ∈ Auth, rq ≤AS subject, uri(object) == URI
        OR uri(object) == dtd(URI)};
    For a in A do
        N = {n | n ∈ T, n ∈ a.object};
        Case a.sign of
            '+': For n in N do n.vclabel[a.type].Allowed.Add(a.subject);
            '-': For n in N do n.vclabel[a.type].Denied.Add(a.subject);
        }
}

void SetLabel()
/* Evaluates the set of authorizations of each type on the node */
{
    For t in [LDH, RDH, L, R, LD, RD, LS, RS]
    {
        s = this.vclabel[t].Denied.Head(); keepSubject = TRUE;
        While s != null do
        {
            s' = this.vclabel[t].Allowed.Head();
            While s' != null AND keepSubject do
            {
                If (s ≤AS s' AND ¬(s' ≤AS s)) /* s' dominated and "most specific takes precedence" */
                then this.vclabel[t].Allowed.Remove(s');
                else If (s' ≤AS s AND ¬(s ≤AS s'))
                then keepSubject = FALSE; s' = this.vclabel[t].Allowed.Next();
            }
            If ¬keepSubject /* s dominated and "most specific takes precedence" */
            then this.vclabel[t].Denied.Remove(s);
            s = this.vclabel[t].Denied.Next();
        }
        If this.vclabel[t].Allowed.IsEmpty()
        then If this.vclabel[t].Denied.IsEmpty()
            then vclabel[t].sign = 'ε';
            else vclabel[t].sign = '-';
        else If this.vclabel[t].Denied.IsEmpty()
            then vclabel[t].sign = '+';
            else vclabel[t].sign = CONFLICTPOLICY; /* '-' for the "denials take precedence" policy */
    }
    For c in this.Children() do c.SetLabel()
}

void GetFinallabel(pveclabel)
{
    this.finlabel = 'ε';
    For t in [LDH, RDH, L, R, LD, RD, LS, RS]
    {
        this.vclabel[t].sign = this.vclabel[t].sign ⊕ pveclabel[t].sign;
        this.finlabel = this.finlabel ⊕ this.vclabel[t].sign;
    }
    For a in this.Attribute() do a.GetFinallabel(this.vclabel);
    For e in this.ElemChildren() do e.GetFinallabel(masklocal(this.vclabel));
}

void Prune()
{
    For c in this.Children() do c.Prune();
    If this.Children() == ∅ and this.finlabel ≠ '+' then remove the current node from T;
}

```

Fig. 6. Compute-view algorithm.

principle. We refer to the combination of these two conflict resolution policies as “*most specific subject/denials take precedence*” principle. The reason for this specific choice is that the two principles so combined naturally cover the intuitive interpretation that one would expect from the specifications [Lunt 1989]. This behavior is realized by method **SetLabel**, applied on all the document nodes in a preorder visit starting from the root r . **SetLabel** combines, for each type t , the two lists $veclabel[t].Allowed$ and $veclabel[t].Denied$ of subjects according to the “*most specific subject/denials take precedence*” principle. Intuitively, a given subject s belonging to list $veclabel[t].Denied$ is compared with each subject s' in the list $veclabel[t].Allowed$. If s is more specific than s' , s' can be removed from $veclabel[t].Allowed$ as it is dominated by s . If, on the contrary, s' is more specific than s , it is s that is dominated and can be removed from the list $veclabel[t].Denied$. When all the subjects appearing in $veclabel[t].Denied$ have been compared with the subjects in $veclabel[t].Allowed$, the content of the two lists is considered: if the two lists are empty, this means that no authorization was originally defined on the node and the value ε is assigned to $veclabel[t].sign$; if the list $veclabel[t].Allowed$ is empty and the list $veclabel[t].Denied$ is not empty, only negative authorizations are applicable on the node and $-$ is assigned to $veclabel[t].sign$; if the list $veclabel[t].Allowed$ is not empty and the list $veclabel[t].Denied$ is empty, only positive authorizations are applicable on the node and $+$ is assigned to $veclabel[t].sign$; finally, if both the lists are not empty, this means that there is a conflict where authorizations with unordered subjects have been defined on the same node and the sign specified by the conflict resolution policy must be assigned to $veclabel[t].sign$ (in our case, $-$). It is important to note, however, that our model can support any of the conflict resolution policies discussed. Indeed, a different policy requires only a change in method **SetLabel**. Also, different policies could be applied to the same server, towards the definition of multiple policy systems [Jajodia et al. 2001]. The only obvious restriction we impose is that no more than one policy apply for each document.

The labels (signs) associated with nodes are then propagated to their subelements and attributes according to the following criteria: (1) authorizations on a node take precedence over those on its ancestors, and (2) authorizations at the instance level, unless declared as soft, take precedence over authorizations at the schema level, unless declared as hard. The complete labeling of the document tree is thus obtained by calling method **GetFinalLabel**($veclabel$) on root node r of T . Method **GetFinalLabel** considers the nodes of T according to a preorder visit of the tree and propagates permissions/denials associated with a node to its descendants. Propagation of the value, for each type t , of $p.veclabel[t].sign$ of a node p parent of a node n is obtained by assigning to $n.veclabel[t].sign$ the value of $p.veclabel[t].sign$ if and only if $n.veclabel[t].sign$ is equal to ε . This propagation can be performed by interpreting the three values $+$, $-$, and ε as values of a three-valued logic. To this end, we first need to map $+$, $-$, ε in the logic. The only condition that such mapping must satisfy is that ε must be mapped to 0 (false). To understand the reason for this, think of false as “no statement” has been made. Signs $+$ and $-$ must then be mapped to the other two values, namely, 1 (true), and $\frac{1}{2}$ (indeterminate); whatever choice would do.


A	B	A \wedge B	A	B	A \vee B	A	\neg A	A	B	A \oplus B
ε	ε	ε	ε	ε	ε	ε	+	ε	ε	ε
ε	-	ε	ε	-	-	-	-	ε	-	-
ε	+	ε	ε	+	+	+	ε	ε	+	+
-	ε	ε	-	ε	-	-	-	-	ε	-
-	-	-	-	-	-	-	-	-	-	-
-	+	-	-	+	+	-	-	-	+	-
+	ε	ε	+	ε	+	+	+	+	ε	+
+	-	-	+	-	+	+	-	+	-	+
+	+	+	+	+	+	+	+	+	+	+

Fig. 7. Truth tables of the propositional connectives \wedge , \vee , and \neg , and operator \oplus .

Here, we map + to 1 and - to $\frac{1}{2}$. It is easy now to see that, with the defined mapping, the propagation is obtained by assigning to $n.\text{veclabel}[t].\text{sign}$ the result of the formula $n.\text{veclabel}[t].\text{sign} \vee (\neg n.\text{veclabel}[t].\text{sign} \wedge p.\text{veclabel}[t].\text{sign})$, where the truth tables of the propositional connectives \vee , \neg , and \wedge coincide with the truth tables defined in the three-valued logic of Lukasiewicz [Rescher 1969] (see Figure 7). We denote such a formula as $n.\text{veclabel}[t].\text{sign} \oplus p.\text{veclabel}[t].\text{sign}$ in the following. The truth table for \oplus is reported in Figure 7. In the case where n is an element, propagation follows the same principle but $n.\text{veclabel}$ is combined with a masked version of the parent array $p.\text{veclabel}$ obtained by means of function **masklocal** that sets to ε the *sign* field of components LDH, L, LD, and LS. The reason for this is that local authorizations applicable to a node p can be propagated only to attributes of p . After this propagation step, according to the defined priorities, **GetFinalLabel** determines the sign *finlabel* that must hold for the specific node n . In particular, the final sign *finlabel* of each node n is determined as the result of operation \oplus between the *sign* field of components of array $n.\text{veclabel}$ considered in their priority order: LDH (local hard), RDH (recursive hard), L (local), R (recursive), LD (local, schema level), RD (recursive, schema level), LS (local soft), and RS (recursive soft).

6.2 Transformation Process

As a result of the labeling process, the value of *finlabel* for each node n contains the sign, if any, reflecting whether the node can be accessed (+) or not (-). The value of *finlabel* is equal to ε in the case where no authorizations have been specified nor can be derived for n . Value ε can be interpreted either as a negation or as a permission, corresponding to the enforcement of the closed and the open policy, respectively [Jajodia et al. 2001]. In the following, we assume the closed policy. Accordingly, the requester is allowed to access all the elements and attributes whose label is positive. To preserve the structure of the document, the portion of the document visible to the requester will also include start and end tags of elements with a negative or undefined label that have a descendant with a positive label. The view of the document can be obtained by pruning from the original document tree all the subtrees containing only nodes labeled negative or undefined. This pruning is enforced by method **Prune** in Figure 6, which executes a postorder visit on the document tree and removes any leaf labeled - or ε . The pruned document may be not valid with respect to the DTD referenced by the original XML document. This may happen, for instance, when required

Fig. 8. Execution steps of the **Compute-view** algorithm.

attributes are deleted. To avoid this problem, a *loosening* transformation is applied to the DTD. Loosening a DTD simply means to define as optional all the elements and attributes marked as required in the original DTD. DTD loosening prevents users from detecting whether information was hidden by the security enforcement or was simply missing in the original document.

Figure 8 summarizes all the execution steps of the **Compute-view** algorithm.

Example 6.1. Consider the set of authorizations defined in Figure 5, and the user-group hierarchy in Figure 3(a). Consider a request to read the XML document in Figure 2 submitted by user Alice connected from host 159.101.80.10 with symbolic name *tweety.cardiology.hospital.com*. According to the authorizations stated at the DTD and instance levels, since Alice is a member of the medical staff of the hospital, she can only access medical information such as the names of patients, their rooms, the drug names, and the corresponding daily administration quantity given to patients. Consider now the same request submitted by user Tom connected from host 159.101.80.5 with symbolic name *hole.admin.hospital.com*. Since Tom is a member of group Administrative, he can access administrative information such as the name, home address, and salary of the medical staff. By contrast, not belonging to the medical staff he cannot access medical information on patients (e.g., illness, type of therapy, and drug name). Figures 9(a) and (b) show the resulting view of the documents returned to Alice and Tom, respectively. These views reflect a general principle according to which each user can access only information needed to complete his or her activity (*need-to-know* principle [Samarati and De Capitani di Vimercati 2001]): Alice's view contains all the information that she might need as a nurse, and Tom's view contains information related to administrative tasks. Neither Alice nor Tom have a full view of the document.

We conclude this section with a mention of the performance characteristics of our system. The two tasks that have the greatest impact on performance are the identification of the authorization objects and the evaluation of node labels.

Authorization objects are identified by XPath expressions, which are evaluated inside the **InitialLabel** procedure. This step dominates the complexity in the system, as XPath is a rich programming language that permits the definition of search expressions requiring exponential time for their evaluation [Mendelzon and Wood 1995]. Several works identify restrictions on XPath that reduce its complexity characteristics [Buneman et al. 2000; Deutsch and Tannen 2001; Mendelzon and Wood 1995], keeping a level of expressivity adequate for most situations.

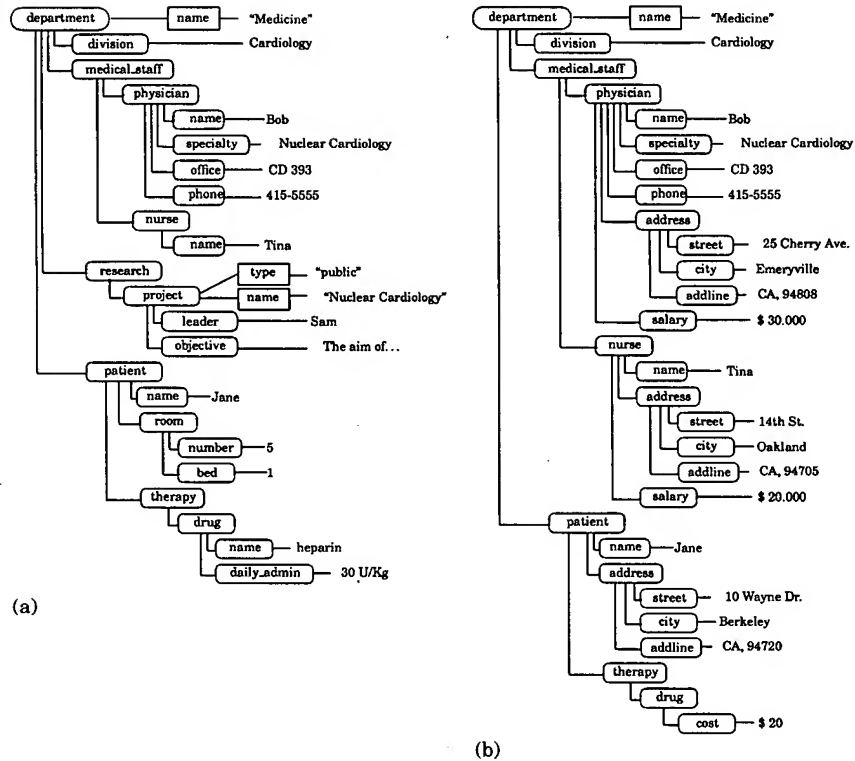


Fig. 9. The view (a) of user Alice and (b) of user Tom.

The evaluation of node labels, although expressive and flexible, bears limited computational cost. This can be noticed by quickly evaluating Algorithm 6. In fact, the tree labeling initialization (**InitialLabel**) is linear in the number of authorizations associated with the document, whereas the label computation (**GetFinalLabel**) and the pruning process (**Prune**) are linear in the number of nodes in the document. The most expensive operation seems therefore the set label computation (**SetLabel**) that for each node in the tree enforces the conflict resolution policy for solving inconsistencies among authorizations of the same type. In principle such an operation could have a worst-case cost, for each node, quadratic in the highest number of authorizations of a given type that are associated with the node (subject comparison for specificity being assumed constant [Raynaud and Thierry 2001]). It is, however, legitimate to assume the number of such authorizations to be very small, and limited by a reasonable constant, making the **SetLabel** method also linear in the number of nodes in the document.

7. SUPPORTING WRITE ACTIONS

In previous sections, we treated only read authorizations. This is justified by practical considerations, as currently XML applications are mostly read-only.

Read authorizations also permit a simpler description of the approach, since the emergence of conflicts is easier to understand if only read privileges are involved. Finally, whereas read operations can be immediately modeled, no consensus has emerged up to now in the research community on a model for XML updates. In this section, we introduce a basic model for XML writes that permits the introduction of write authorizations. Richer models have already been defined for the representation of XML writes (e.g., Abiteboul et al [1999], Goldman et al [1999], and Liefke and Davidson [2000]) and write authorizations could be specialized for a specific model, with direct support for the complex operations that the model offers (e.g., movement of nodes in the Lorel model [Goldman et al 1999] or the merging of XML trees in the WHAX model [Liefke and Davidson 2000]), but these are specific customizations that we do not treat here. Like read authorizations, write authorizations can be local or recursive, hard or soft, and can be specified on elements/attributes within either single XML documents or DTDs. The semantics of local and recursive authorizations remains unchanged; local authorizations on elements/attributes apply only to the considered elements/attributes and recursive authorizations apply also to their subelements. Write authorizations specified at the DTD level apply to all the DTD instances and write authorizations specified at the instance level apply only to the document on which they are defined. Conflict resolution is applied in the same way as for read operations, with a complete separation among authorizations on different actions.⁶

We define write operations with a basic model where only operations on single nodes are considered. The operations on the node can be insert the node, delete the node, and update the node (i.e., change the value for an attribute or a change of the text for an element). The three operations correspond to distinct write privileges: insert, delete, and update. Note that insert privileges allow the insertion of new elements and attributes in a document, which, although not existing before the insert operation itself, can be specified with reference to the schema (DTD). For instance, an insert authorization on element *therapy* of a document allows the insertion of a new *therapy* within the document. The consideration of the three write privileges above offers most of the services required for an access control mechanism on write operations. Being at low level, it is also compatible with many models for the representation of XML operations (e.g., graph-based [Goldman et al. 1999], object-based [Abiteboul et al. 1999], tree-based [Liefke and Davidson 2000], or object-relational [Oracle Corp. 2000]), which can all be mapped in terms of simple operations.

Insert operations are evaluated by executing the labeling process on the document with the new node inserted. If the labeling produces a positive label on the new node the insert completes successfully, otherwise the document remains unchanged. Note that possible conditions associated with authorizations can be exploited to specify restrictions on the document where the insertion can be

⁶The model could be extended to the support of action hierarchies in the form of authorization implication (e.g., a write access implies a read access) or precondition (e.g., a write access on an element requires at least a read access on the element's ancestors) [Sandhu and Samarati 1997].

executed as well as on the values that can be inserted. For instance, authorization `((PhyC,*,*.hospital.com),/department[./@name='medicine']/patient//therapy[./cost < 10,000],insert,+,R)` states that members of group PhyC can insert a new therapy for a patient in the Medicine department, and that the privilege is limited to insertion of therapies whose cost does not exceed \$10,000.

Deletion of a node is permitted only if the labeling of the document produces a final positive label for the node to be deleted. If so, the node is eliminated and will not be present in the new document. For instance, authorization `((Administrative,159.101.80.77,secws.hospital.com),/department/medical_staff//bonus,delete,+,L)` states that users of the Administrative group connected from the specified location have the privilege to delete element bonus of the members of the medical staff.

Updates are evaluated by executing the labeling process on both the existing document and on the document that results from the execution of the update. If the final label associated with the node being updated is positive in both versions the operation is completed successfully; otherwise it is rejected. For instance, authorization `((NurseC,*,*.hospital.com),/department[./@name='medicine']/patient/room/bed[number(value())>=100 and number(value())<=150],update,+,L)` specifies that members of group NurseC can update element bed of patients, only if the bed belongs to the block of beds 100 through 150. They are thus responsible for updating the distribution of patients in the block, but they are not allowed to transfer patients across blocks. From the example it is then clear that the reason for considering both the original and the updated version of the document is to require the satisfaction of the condition in the authorizations in both the old and the new document's state [Atzeni et al. 1999].⁷

When the XML document is modified, the system must also check the correctness of the document with respect to the DTD and, if the document is not valid, the write must not be accepted. Incidentally, if the document is characterized by a loosened DTD, this check can have an impact only on insertions and updates: insertions can introduce elements with a tag incompatible with the DTD, and updates can specify a value in contrast to what the DTD specifies for the node; deletions are always accepted. With a generic DTD and possibly with semantically richer specifications, such as those proposed by XML schema [Thompson et al. 2001; Biron and Malhotra 2001], the constraints that must be verified on the document can become quite complex and checks are required for every action.

The model above considers write operations on single nodes. However, write requests often refer to sets of nodes (typically a subtree in the document). In this case it may be convenient to introduce a transactional mechanism based on the “deferral” of controls, analogous to the SQL command set constraints

⁷We note that alternative approaches could interpret the conditions in the authorizations only as conditions on the values being updated, in which case only the original document should be labeled (as for the delete operation) or only as conditions on the new values being introduced, in which case only the new document should be labeled (as for the insert operation).

deferred that relational systems offer for the management of constraints. The idea is that writes are collected in an atomic sequence, and all the checks on the correctness of the updates are deferred at the end of the sequence, when each single write is considered for permissions and correctness. If a single write is not permitted, the sequence is invalid and the XML data are rolled back to their original state. The transactional mechanism also allows the support of coordinated write operations that are individually incorrect but globally correct. We can consider this example. Suppose a user is given the authorization to manage the description of patients that are in the Medicine department and are in a critical condition. If a control is executed on every single write of a node, the user may be forbidden to insert new patients: the insertion of a new patient node is not allowed because the patient node has no other node and the user is not permitted to modify patients who are not in a critical condition. By contrast, deferring the control at the end of the update sequence, the write operation can be successfully completed.

8. SOFTWARE ARCHITECTURE

We briefly present the architecture of the *XML access control processor* (ACP), the component that wraps up entirely the computation of access permissions to individual elements and performs the corresponding transformation on the XML document.

8.1 Architectural Requirements

From the architectural point of view, a first point to note is that our design is fully server side. The choice between *server-side* and *client-side* processing is a typical one when XML is used in the Web context [Park et al. 2001]. As the access control process should clearly be trusted to operate correctly and restrict each requestor to the data he is entitled to access, it is highly preferable to use server-side processing; otherwise protected information should be transmitted to the client and a complex infrastructure should be implemented to guarantee clients' trust in properly enforcing the access restrictions. The satisfaction of the requirements specific to the protection of information is not sufficient to guarantee a successful adoption of this technology in real applications. Indeed, the following requirements must also be considered.

- *Seamless integration*: XML access control should be provided as seamlessly as possible, without interfering with the operation of other presentation or data-processing services. Moreover, the access control service should be introduced on existing servers with minimal interruption of their operation.
- *Quality of service*: The emergence of the World Wide Web as a mainstream technology has highlighted the problem of providing a high quality of service (QoS) to application users. This factor alone should caution about the risk of substantially increasing the processing load of Web servers.

With respect to the first requirement, the key technology used for the integration of access control with other services is the document object management (DOM) specification [World Wide Web Consortium (W3C) 1998], an API defined

Table III. Advantages and Disadvantages of Concurrency Control Techniques

Technique	Advantages	Disadvantages
Single threaded	No context switch overhead	Not scalable on multiprocessors
Process-per-request	Portability	Resource intensive
Process pool	No process creation costs	Not available on every OS
Thread-per-request	Speed	Requires mutual exclusion
Thread pool	Speed	Mutual exclusion on some OS

by the W3C to process XML information. Several systems implement the DOM interface, in various languages; each of these systems offers services for the bidirectional transformation between the textual representation of an XML file and an internal proprietary representation, on which the methods of the API operate. The use of the DOM API offers great potential in the integration of different components managing XML information, because each component can be designed as a DOM transformer independent of the others. For instance, our access control processor can operate on the DOM representation produced by a cache manager and the result of its simplification may be passed to an XML query engine computing a new document.

The system has been designed following the principles of object-orientation and is based on the specification of a set of classes. The classes have been defined in an abstract way using the interface definition language (IDL) [Mowbray and Malveau 1997]. The classes are organized into families. The first family is an extension of the DOM class hierarchy and enriches the description of the nodes of a document with the required security attributes. This extension makes use of the inheritance mechanism that permits an immediate integration with existing DOM implementations. The second family of classes is strictly related to the processing of the access control model and describes all the concepts that are part of the model, such as authorization signs, subjects, path expressions, and complete authorizations.

The quality of service issue centers on performance. Several techniques are available to implement Web-based, high-concurrency systems, each having its positive and negative aspects, as illustrated in Table III. Multithreaded designs are currently the preferred choice for Web servers, as the cost of spawning a thread is usually much lower than that of a process. This is also our design choice for our ACP implementation. Beside being usable in a single-thread execution, our processor can be easily interfaced to a *Dispatcher* registered with an *Event Handler*. Dispatcher-based multithreading can be managed synchronously, according to the Reactor/Proactor design pattern [Lavender and Schmidt 1995], or asynchronously, as in the Active Object pattern. We adopted the former choice, as it further facilitates the integration of our XML access control code in the framework of existing general-purpose server-side transformers based on the same design pattern (such as Cocoon [Apache Software Foundation 2000]). To obtain a multithreaded system, the classes must be implemented in a thread-safe way, making all the parameters of each request isolated in the context of a specific method invocation.

An architectural solution that further enhances the performance is the separation of the threads managing the user hierarchy. The services that compute

<pre> <ELEMENT set_of_authorizations (authorization)+> <ELEMENT authorization (subject,object,action,sign,type)> <ELEMENT subject (#PCDATA)> <ELEMENT object (#PCDATA)> <ELEMENT action empty> <ELEMENT sign empty> <ELEMENT type empty> </pre>	<pre> <ATTLIST set_of_authorizations about CDATA #REQUIRED> <ATTLIST authorization type value (L RLD RDH L RS L RD) #REQUIRED> sign value (+ -) #REQUIRED> <ATTLIST action value (read) #REQUIRED> </pre>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Fig. 10. XAS syntax.

if an authorization is applicable to a given user need to evaluate if a user is a member, directly or indirectly, of the group specified in the authorization's subject. The efficiency of this computation can be greatly increased by building auxiliary indices [Agrawal et al. 1989] that synthetically describe the membership in groups. These structures are expensive to build and describe relatively static information, thus it is best to build them once in a thread that always stays active and offers its services to the threads managing the document transformation.

Explicit synchronization mechanisms must be used to ensure that conflicting write requests for shared resources are correctly managed. In the ACP object, the only shared resource where writes can occur is the user hierarchy, but requests for user/group changes are infrequent, thus the synchronization does not have an impact on system performance.

8.2 Execution Phases

We are now ready to describe how the ACP works. Our processor takes as input a valid XML document requested by the user, together with its XML Access Sheet (XAS) listing the associated access authorizations at the instance level. The processor operation also involves the document's DTD and the associated XAS specifying schema-level authorizations. The processor output is a valid XML document including only the information the user is allowed to access. To provide a uniform representation of XASs and other XML-based information, the syntax of XASs is given by the XML DTD depicted in Figure 10. The XASs associated with an XML document and its DTD are located by relying on the abstract nature of the XML XLink specification [DeRose et al. 2001] to define out-of-line links that reside outside the documents they connect, making links themselves a viable and manageable resource. The repertoire of out-of-line links defining access control mappings is itself an XML document, easily managed and updated by the system manager; nonetheless it is easily secured by standard file-system level access control.

Our security processor computes an online transformation on XML documents. Its execution cycle, illustrated in Figure 11, consists of the following basic steps.

1. *Parsing.* The parsing step consists of the syntax check of the requested document with respect to the associated DTD and its compilation to obtain an object-oriented document graph according to the DOM format. Since the parsing is performed externally when the ACP is used as a transformer in the framework of a complete architecture complying with the well-known Pipes and Filters design pattern [Buschmann et al. 1996], here we do not deal with parsing issues in detail.

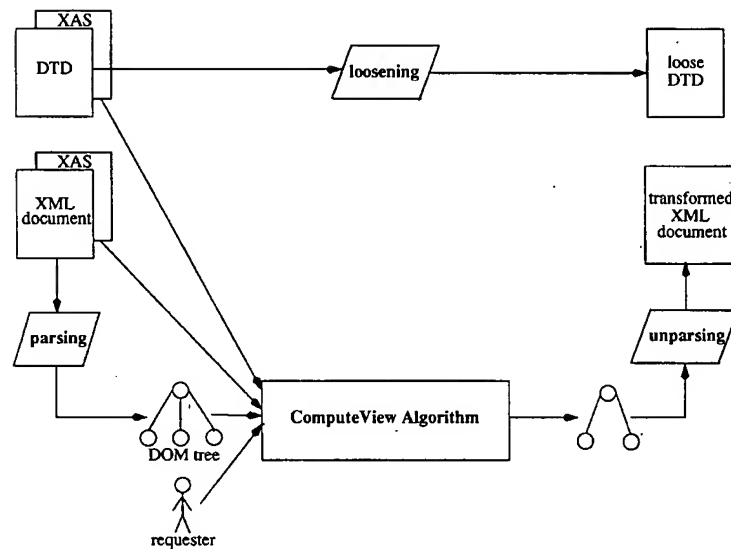


Fig. 11. Execution steps of the security processor.

2. *Compute view.* The compute view step determines the requester's view, by applying the algorithm presented in Section 6 and according to the authorizations listed in the XASs associated with the document and its DTD. As already discussed, the resulting view preserves the validity of the document with respect to the loosened version of its original DTD.
3. *Unparsing.* Finally, the third step is the generation of a valid XML document in text format, simply by unparsing (again, by means of a standard component) the DOM tree computed by the previous step. Once again, this step is performed externally when the ACP is executed as a transformer in the framework of a Pipes and Filters system.

8.3 The Java Implementation

We designed the prototype of the access control model (<http://seclab.dti.unimi.it/~xml-sec>) in Java, using the services of a Java implementation of the DOM API (we used the IBM's XML4J processor [AlphaWorks 2001], which has now evolved into the Apache's Xalan tool [Foundation 2001]). This choice had several consequences on the behavior of our prototype. Here we make a few observations that strictly depend on this choice.

First, it must be noted that no Java-based design of multithreading components has full control of thread management: when running on an operating system that supports threads, the Java virtual machine (JVM) automatically maps Java threads to native threads [Lea 1996], whereas when no native thread support is available, the JVM has to emulate threads. In the latter case, the emulation technique chosen by the JVM implementors can have a significant impact on performance.

Another interesting point to discuss is the mechanism used for the integration of the ACP with the HTTP server. A simple technique consists of the adoption of the common gateway interface (CGI), which allows a Web server to run a generic executable when managing a request; this solution is available independently from the language used to implement the IDL specification. A Java implementation offers specific solutions and indeed our prototype is invoked by the HTTP server using Java servlets. The servlet specification [Sun Microsystems 1999b] defines a protocol for the exchange of information between the server and the JVM, where a request to the server for a resource identified by a URL generates the invocation of the JVM from the server, passing all the parameters' part of the request. Overall, servlets constitute an interesting solution for the implementation of Web services, offering a higher-level interface than CGI, with a relatively easy integration between the HTTP server and the Java environment. Other solutions, specific for Java, are available for the prototype implementation. In particular, for the next version of the prototype, we are investigating the use of Java server pages (JSP) [Sun Microsystems 1999a], a technology built upon servlets offering template-based invocation of Java services.

9. CONCLUSIONS

The definition of an authorization mechanism for protecting data offered on Web sites is an important research direction and a practical pressing need. Existing proposals, specifying protection requirements at the file-system level or with reference to the HTML constructs, turn out to be very limited. By exploiting the opportunities offered by XML, we have defined an access control model for restricting access to Web documents that takes into consideration the semistructured organization of data and their semantics. The result is an access control system that, although powerful and able to represent different protection requirements easily, proves simple and of easy integration with existing applications. Our proposal leaves space for further work. Issues to be investigated include: the consideration of requests in the form of generic queries, extension of the model to role- and credential-based authorizations, and the investigation of performance optimizations.

REFERENCES

- ABITEBOUL, S., AMANN, B., CLUET, S., EYAL, A., MIGNET, L., AND MILO, T. 1999. Active views for electronic commerce. In *Proceedings of the 25th International Conference on Very Large Data Bases* (Edinburgh, Scotland, Sept.), 138–149.
- ADLER, S., BERGLUND, A., CARUSO, J., DEACH, S., GRAHAM, T., GROSSO, P., GUTENTAG, E., MILOWSKI, A., PARNELL, S., RICHMAN, J., AND ZILLES, S. 2001. Extensible stylesheet language (XSL) version 1.0. World Wide Web Consortium (W3C). Available at <http://www.w3.org/TR/xsl>.
- AGRAWAL, R., BORGIDA, A., AND JAGADISH, H. V. 1989. Efficient management of transitive relationships in large data and knowledge bases. In *Proceedings of the ACM SIGMOD'89 Conference* (Portland, Or., May–June), ACM, New York, 253–262.
- ALPHAWORKS. 2001. XML security suite. Available at <http://www.alphaWorks.ibm.com/tech/xmlsecuritysuite>.
- APACHE SOFTWARE FOUNDATION. 2000. Cocoon, a Java publishing framework. Available at <http://xml.apache.org/cocoon>.

- APACHE SOFTWARE FOUNDATION. 2001. Xalan-J version 2.2.d14. Available at <http://xml.apache.org/xalan-j/>.
- ATZENI, P., CERI, S., PARABOSCHI, S., AND TORLONE, R. 1999. *Database Systems—Concepts, Languages and Architectures*. McGraw-Hill, New York.
- BERNERS-LEE, T., FIELDING, R., IRVINE, U., AND MASINTER, L. 1998. Uniform resource identifiers (URI): Generic syntax. Available at <http://www.isi.edu/in-notes/rfc2396.txt>.
- BERTINO, E., CASTANO, S., AND FERRARI, E. 2001. Securing XML documents with author-x. *IEEE Internet Comput.* 5, 3 (May/June), 21–31.
- BERTINO, E., CASTANO, S., FERRARI, E., AND MESITI, M. 2000. Specifying and enforcing access control policies for XML document sources. *World Wide Web J.* 3, 3.
- BIRON, P. AND MALHOTRA, A. 2001. XML schema part 2: Datatypes. World Wide Web Consortium (W3C). Available at <http://www.w3.org/TR/xmlschema-2>.
- BRAY, T., HOLLANDER, D., AND LAYMAN, A. 1999. Namespaces in XML. World Wide Web Consortium (W3C). Available at <http://www.w3.org/TR/REC-xml-names>.
- BRAY, T., PAOLI, J., SPERBERG-McQUEEN, C., AND MALER, E. 2000. Extensible markup language (XML) 1.0 (second edition). World Wide Web Consortium (W3C). Available at <http://www.w3.org/TR/REC-xml>.
- BRICKLEY, D. AND GUHA, R. 2000. Resource description framework (RDF) schema specification 1.0. World Wide Web Consortium (W3C). Available at <http://www.w3.org/TR/rdf-schema>.
- BUNEMAN, P., FAN, W., AND WEINSTEIN, S. 2000. Path constraints in semistructured databases. *JCSS* 61, 2, 146–193.
- BUSCHMANN, F., MEUNIER, R., ROHNERT, H., SOMMERLAD, P., AND STAL, M. 1996. *Pattern-Oriented Software Architecture—A System of Patterns*. Wiley, New York.
- CHECKFREE CORP. 2001. Open financial exchange specification 2.0.1. CheckFree Corp. Available at <http://www.ofx.net/>.
- CONTENT GUARD. 2001. eXtensible rights markup language (XrML) 2.0. ContentGuard. Available at <http://www.xrml.org/>.
- DAMIANI, E., DE CAPITANI DI VIMERCATI, S., PARABOSCHI, S., AND SAMARATI, P. 2000a. Design and implementation of an access control processor for XML documents. *Comput. Netw.* 33, 1–6 (June), 59–75.
- DAMIANI, E., DE CAPITANI DI VIMERCATI, S., PARABOSCHI, S., AND SAMARATI, P. 2000b. Securing XML documents. In *Proceedings of EDBT 2000* (Konstanz, Germany, March), *Lecture Notes in Computer Science*, vol. 1777, Springer, New York, 121–135.
- DEROSE, S., MALER, E., AND ORCHARD, D. 2001. XML linking language (XLink) version 1.0. World Wide Web Consortium (W3C). Available at <http://www.w3.org/TR/xlink>.
- DEUTSCH, A. AND TANNEN, V. 2001. Containment and integrity constraints for xpath. In *Proceedings of the Eighth International Workshop on Knowledge Representation Meets Databases* (Rome, Sept.).
- DEVANBU, P., GERTZ, M., KWONG, A., MARTEL, C., NUCKOLLS, G., AND STUBBLEBINE, S. 2001. Flexible authentication of XML documents. In *Proceedings of the Eighth ACM Conference on Computer and Communications Security* (Philadelphia, Nov.).
- EASTLAKE, D., REAGLE, J., AND SOLO, D. 2001. XML-signature syntax and processing. Available at <ftp://ftp.rfc-editor.org/in-notes/rfc3075.txt>.
- ELLERMAN, C. 1997. Channel definition format (CDF). Available at <http://www.w3.org/TR/NOTE-CDFsubmit.html>.
- FERNANDEZ, E., GUEDES, E., AND SONG, H. 1994. A model of evaluation and administration of security in object-oriented databases. *IEEE Trans. Knowl. Data Eng.* 6, 2 (April), 275–292.
- GABILLON, A. AND BRUNO, E. 2001. Regulating access to XML documents. In *Proceedings of the Fifteenth Annual IFIP WG 11.3 Conference on Database Security* (Niagara on the Lake, Ontario, July).
- GOLDMAN, R., MCHUGH, J., AND WIDOM, J. 1999. From semistructured data to XML: Migrating the lore data model and query language. In *Proceedings of the Second International Workshop on the Web and Databases (WebDB '99)* (Philadelphia, June).
- JAJODIA, S., SAMARATI, P., SAPINO, M., AND SUBRAHMANYAN, V. 2001. Flexible support for multiple access control policies. *ACM Trans. on Database Syst.* 26, 2 (June), 214–260.

- JONSCHER, D., MOFFETT, J., AND DITTRICH, K. 1994. Complex objects or: The striving for complexity is ruling our world. In *Database Security, VII: Status and Prospects*, T. Keefe and C. Landwehr, Eds., Elsevier Science, Amsterdam, 19–37.
- KUDOH, M., HIRAYAMA, Y., HADA, S., AND VOLLSCHWITZ, A. 2000. Access control specification based on policy evaluation and enforcement model and specification language. In *Symposium on Cryptography and Information Security, SCIS'2000*.
- LAVENDER, R. G. AND SCHMIDT, D. 1995. Reactor: An object behavioral pattern for concurrent programming. In *Pattern Languages of Program Design 2*. J. Vlissides, D. Coplien, and M. Kerth, Eds., Addison-Wesley, Reading, Mass.
- LEA, D. 1996. *Concurrent Programming in Java*. Addison-Wesley, Reading, Mass.
- LIEPKE, H. AND DAVIDSON, S. 2000. View maintenance for hierarchical semistructured data. In *Proceedings of Data Warehousing and Knowledge Discovery (DaWaK 2000)* (London Greenwich, UK, Sept.).
- LUNT, T. 1989. Access control policies for database systems. In *Database Security, II: Status and Prospects*, C. Landwehr, Ed., North-Holland, Amsterdam, 41–52.
- MENDELZON, A. AND WOOD, P. 1995. Finding regular simple paths in graph databases. *SIAM J. Comput.* 24, 6, 1235–1258.
- MOWBRAY, T. AND MALVEAU, R. 1997. *CORBA Design Patterns*. Wiley, New York.
- ORACLE CORP. 2000. Oracle and XML. Available at <http://www.oracle.com/xml>.
- PARK, J., SANDHU, R., AND AHN, G. 2001. Role-based access control on the web. *ACM Trans. Inf. Syst. Sec.* 4, 1 (Feb.), 37–71.
- RABITTI, F., BERTINO, E., KIM, W., AND WOELK, D. 1991. A model of authorization for next-generation database systems. *ACM Trans. Database Syst.* 16, 1 (March), 89–131.
- RAYNAUD, O. AND THIERRY, E. 2001. A quasi optimal bit-vector encoding of tree hierarchies. Application to efficient type inclusion tests. In *Proceedings of the Fifteenth European Conference on Object-Oriented Programming (ECOOP)* (Budapest, June).
- REAGLE, J. AND CRANOR, L. 1999. The platform for privacy preferences. *Commun. ACM* 42, 2 (Feb.), 48–55.
- RESCHER, N. 1969. *Many Valued Logics*. Mc Graw-Hill, New York.
- RESCORLA, E. AND SCHIFFMAN, A. 1999. The secure hypertext transfer protocol. Available at <http://www.ietf.org/rfc/rfc2660.txt>.
- SAMARATI, P. AND DE CAPITANI DI VIMERCATI, S. 2001. Access control: Policies, models, and mechanisms. In *Foundations of Security Analysis and Design*, R. Focardi and R. Gorrieri, Eds., *Lecture Notes in Computer Science*, vol. 2171. Springer-Verlag, New York.
- SAMARATI, P., BERTINO, E., AND JAJODIA, S. 1996. An authorization model for a distributed hypertext system. *IEEE Trans. Knowl. Data Eng.* 8, 4 (August), 555–562.
- SANDHU, R. AND SAMARATI, P. 1997. Authentication, access control and intrusion detection. In *CRC Handbook of Computer Science and Engineering*, A. Tucker, Ed., CRC Press, Boca Raton, Fla., 1929–1948.
- SUN MICROSYSTEMS. 1999a. Java server pages (JSP) specification, release 1.1. Available at <http://www.javasoft.com/>.
- SUN MICROSYSTEMS. 1999b. Java servlet API specification, release 2.2. Available at <http://www.java soft.com/>.
- THOMPSON, H., BEECH, D., MALONEY, M., AND MENDELSON, N. 2001. XML schema part 1: Structures. World Wide Web Consortium (W3C). Available at <http://www.w3.org/TR/xmlschema-1>.
- VAN HOFF, A., PARTOVI, H., AND THAI, T. 1997. The open software description format (OSD). Available at <http://www.w3.org/TR/NOTE-OSD.html>.
- WOO, T. AND LAM, S. 1993. Authorizations in distributed systems: A new approach. *J. Comput. Sec.* 2, 2,3 (Feb.), 107–136.
- WORLD WIDE WEB CONSORTIUM (W3C). 1998. Document object model (DOM) level 1 specification version 1.0. World Wide Web Consortium (W3C). Available at <http://www.w3.org/TR/REC-DOM-Level-1>.
- WORLD WIDE WEB CONSORTIUM (W3C). 2001. XML path language (XPath) 2.0. World Wide Web Consortium (W3C). Available at <http://www.w3.org/TR/xpath20>.

Received June 2000; revised January 2002; accepted January 2002